arXiv:1702.07815v1 [cs.DS] 25 Feb 2017

Subquadratic Algorithms for the Diameter and the Sum of Pairwise Distances in Planar Graphs*

Sergio Cabello†

February 28, 2017

Abstract

We show how to compute for n-vertex planar graphs in $O(n^{11/6}\operatorname{polylog}(n))$ expected time the diameter and the sum of the pairwise distances. The algorithms work for directed graphs with real weights and no negative cycles. In $O(n^{15/8}\operatorname{polylog}(n))$ expected time we can also compute the number of pairs of vertices at distance smaller than a given threshold. These are the first algorithms for these problems using time $O(n^c)$ for some constant $c < 2$, even when restricted to undirected, unweighted planar graphs.

Keywords: planar graph, diameter, Wiener index, distances in graphs, distance counting, Voronoi diagram.

1 Introduction

Let G be a directed graph with n vertices and abstract arc-lengths $\lambda\colon E(G) \to \mathbb{R}$. The length of a walk in G is the sum of the arc-lengths along the walk. We assume that G has no cycle of negative length. The ***distance*** between two vertices x and y of G, denoted by $d_G(x,y)$, is the minimum length over all paths in G from x to y. While it is common to use the term distance, this is not necessarily a metric. This scenario is an extension of the more common case where the graph G is undirected and the lengths are positive. In that case $d_G(\cdot,\cdot)$ is indeed a metric.

In this paper we are interested in computing basic information about the distances between vertices in the graph G. The ***diameter*** of G is

$$\operatorname{diam}(G) := \max\{d_G(x,y) \mid x,y \in V(G)\},$$

the ***sum of the pairwise distances*** of G is

$$\operatorname{sum}(G) := \sum_{(x,y)\in(V(G))^2} d_G(x,y),$$

and, for any $\delta \in \mathbb{R}$, the ***distance counter*** of G is

$$\operatorname{count}_{\leq}(G,\delta) := |\{(x,y) \in (V(G))^2 \mid d_G(x,y) \leq \delta\}|.$$

For undirected graphs, the value $\operatorname{sum}(G)$ is essentially equivalent to the average distance in the graph and the so-called *Wiener index*. The Wiener index is a basic topological index used in mathematical chemistry with thousands of publications.

*A preliminary version of this work was presented at SODA 2017 [6].

†Department of Mathematics, IMFM, and Department of Mathematics, FMF, University of Ljubljana, Slovenia.
Supported by the Slovenian Research Agency, program P1-0297. Email address: sergio.cabello@fmf.uni-lj.si

Computing the diameter, the sum of the pairwise distances, or the distance counter of a graph is a fundamental problem in graph algorithms. The obvious way to compute them is via solving the all-pairs shortest path problem (APSP) explicitly and then extract the relevant information. A key question is whether one can avoid the explicit computation of all the pairwise distances.

Roditty and Vassilevska Williams [39] show that, for arbitrary graphs with n vertices and $O(n)$ edges, one cannot compute the diameter in $O(n^{2-\delta_0})$ time, for some constant $\delta_0 > 0$, unless the strong exponential time hypothesis (SETH) fails. In fact, their proof shows that for undirected, unweighted graphs we cannot distinguish in $O(n^{2-\delta_0})$ time between sparse graphs that have diameter 2 or 3, assuming the SETH. This implies the same conditional lower bound for computing the sum of the pairwise distances or the distance counter in sparse graphs. Indeed, an unweighted graph G of n vertices has diameter 2 if and only if

$$\text{sum}(G) \;=\; \sum_{x\in V(G)} \Big(\deg_G(x) + 2(n-1-\deg_G(x))\Big) \;=\; 2n(n-1) - 2|E(G)|.$$

Similarly such a graph G has $\text{count}_{\leq}(G,2) = n^2$ if and only if G has diameter 2. Thus, if we could compute the sum of pairwise distances or the distance counter for sparse graphs in $O(n^{2-\delta_0})$ time, we could also distinguish in the same time whether the graph has diameter 2 or larger, and the SETH fails.

Given such conditional lower bounds , it is natural to shift the interest towards identifying families of sparse graphs where one can compute the diameter or the sum of pairwise distances in truly subquadratic time. Here we provide subquadratic algorithms for directed, planar graphs with no negative cycles. More precisely, we show that the diameter and the sum of the pairwise distances can be computed in $O(n^{11/6}\,\text{polylog}(n))$ expected time, while the distance counter can be computed in $O(n^{15/8}\,\text{polylog}(n))$ expected time. There are efficient algorithms for computing all the distances in a planar graph [17] or a specified subset of the distances; see for example [5, 34]. However, none of these tools seem fruitful for computing our statistics in subquadratic time.

Note that our algorithms are the first algorithms using time $O(n^c)$ for some constant $c < 2$, even when restricted to undirected, unweighted planar graphs.

Related work. For graphs of bounded treewidth one can compute the diameter and the sum of pairwise distances in near-linear time [1, 8]. The distance counter for graphs of bounded treewidth can be handled using the same techniques. Recently, Husfeldt [20] has looked at the problem of computing the diameter for undirected, unweighted graphs parameterized by the treewidth and the diameter.

For planar graphs, Wulff-Nilsen [44] gives an algorithm to compute the diameter and the sum of pairwise distances in unweighted, undirected planar graphs in $O(n^2 \log\log n/\log n)$ time, which is slightly subquadratic. Wulff-Nilsen [45] extends the result to weighted directed planar graphs with a time bound of $O(n^2(\log\log n)^4/\log n)$. Note that the running time of these algorithms is not of the type $O(n^c)$ for any constant $c < 2$.

Researchers have also looked into near-optimal approximations. In particular, Weimann and Yuster [43] provide a $(1+\varepsilon)$-approximation to the diameter of undirected planar graphs in $O((n/\varepsilon^4)\,\text{polylog}(n) + 2^{O(1/\varepsilon)}n)$ time. As it was mentioned in [18], a near-linear time randomized $(1+\varepsilon)$-approximation to the sum of pairwise distances in undirected planar graphs can be obtained using random sampling and an oracle for $(1+\varepsilon)$-approximate distances [22, 42]. See [21] for the average distance in arbitrary discrete metric spaces.

Our approach. Let us describe the high-level idea of our approach. The main new ingredient is the use of additively-weighted Voronoi diagrams in pieces of the graph: we make a quite expensive preprocessing step in each piece that permits the efficient computation of such Voronoi diagrams in each piece for several different weights.

To be more precise, let G be a planar graph with n vertices. We first compute an r-division: this is a decomposition of G into $O(n/r)$ pieces, each of them with $O(r)$ vertices and $O(\sqrt{r})$ boundary vertices. This means that all the interaction between a piece P and the complement goes through the $O(\sqrt{r})$ boundary vertices of P.

Consider a piece P and a vertex x outside P. We would like to break P into regions according to the boundary vertex of P that is used in the shortest path from x. This can be modeled as an additively-weighted Voronoi diagram in the piece: each boundary vertex is a weighted site whose weight equals the distance from x. Thus, we have to compute several such Voronoi diagrams for each piece.

Assuming that a piece is embedded, one can treat such a Voronoi diagram as an abstract Voronoi diagram an encode it using the dual graph. In particular, a bisector corresponds to a cycle in the dual graph. We can precompute all possible Voronoi diagrams for $O(1)$ sites, and that information suffices to compute the Voronoi diagram using a randomized incremental construction. Once we have the Voronoi diagram, encoded as a subgraph of the dual graph, we have to extract the information from each Voronoi region. Although this is the general idea, several technical details appear. For example, the technology of abstract Voronoi diagrams can be used only when the sites are cofacial.

We remark that our algorithms actually compute information for the distances from each vertex x of G separately. Thus, for each vertex x we compute the furthest vertex from x, the sum of the distances from x to all vertices, and the number of vertices at distance at most δ from x, for a given $\delta \in \mathbb{R}$. Our main result, Theorem 27, is stated to make this clear.

Assumptions. We will assume that the distance between each pair of vertices is distinct and there is a unique shortest path between each pair of vertices. This can be enforced with high probability using infinitesimal perturbations or deterministically using lexicographic comparison; see for example the discussion in [7]. Since we are aiming to a randomized algorithm and our running times are barely subquadratic, the actual method that is used is not very relevant.

Roadmap. We assume that the reader is familiar with planar graphs. In the next section we explain the notation and some basic background. In Section 3 we explain how to extract information about the vertices contained in a dual cycle. In Section 4 we explain the concept of abstract Voronoi diagrams. In Section 5 we deal with different definitions of Voronoi diagrams in plane graphs and show that they are equivalent. In Section 6 we discuss the algorithmic aspects of computing Voronoi diagrams. In particular, we make an expensive preprocessing to be able to produce Voronoi diagrams faster. In Section 7 we give the data structure that will be used for each piece of an r-division. In Section 8 we give the final algorithms for planar graphs. We conclude with a discussion.

2 Notation and preliminaries

For running times, we use the notation $\tilde{O}(\cdot)$ when we skip polylogarithmic factors in any of the parameters that appears in the statement. For example, if n appears in the discussion, $\tilde{O}(mr)$ means $O(mr \log^c(mnr))$ for some constant c.

For each natural number n, we use the notation $[n] = \{1, \ldots, n\}$. For each set $A \subset \mathbb{R}^2$, we use $\overline{A}$ for its closure and A° for its interior.

Graphs. Graphs considered in this paper are directed. We use $V(G)$ and $E(G)$ for the vertex and the arc set of a graph G, respectively. We use the notation $x \to y$ or e to denote arcs. The tail of an arc $x \to y$ is x, and y is the head. We use e^R for the reversal of the arc e. In some cases we may

have parallel arcs. It should be clear from the context which arc we are referring to. When the orientation of the arc $x \to y$ is not relevant, we may use xy and refer to it as an (undirected) edge.

A closed walk in G is a sequence $e_0, \ldots, e_{k-1}$ of arcs with the property that the tail of e_i is the head of e_{i-1} for all $i \in [k]$ (indices modulo k). Sometimes a closed walk is given as a sequence of vertices. This uniquely defines the closed walk if there are no parallel edges. A cycle is a closed walk that does not repeat any vertex. In particular, a cycle cannot repeat any arcs. We make it clear that the walk $x \to y, y \to x$ is a cycle.

Planarity. A ***plane graph*** is a planar graph together with a fixed embedding. The arcs e and e^R are assumed to be embedded as a single curve with opposite orientations. In the arguments we will use the geometry of the embedding and the plane quite often. For example, we will talk about the faces enclosed by a cycle of the graph. However, all the computations can be done assuming a combinatorial embedding, described as the circular order of the edges incident to each vertex.

Let G^* be the dual graph of a plane graph G. We may consider G^* with oriented arcs or with edges, depending on the context. We keep in G^* any parallel edges that may occur. When G is 2-connected, the graph G^* has no loops. For each vertex v and edge e of G, we use v^* and e^* to denote their dual counterparts, respectively. For any set of edges $A \subseteq E(G)$, we use the notation $A^* = \{e^* \mid e \in A\}$.

We assume natural embeddings of G and G^* where each dual edge e^* of G^* crosses G exactly once and does so at e. There are no other types of intersections between G and G^*. See Figure 1 for an example. If we would prefer to work with an actual embedding and coordinates, instead of a combinatorial embedding, we could do so. To achieve this, for each edge e of G, we subdivide e and e^* with a common vertex v_e. Then we obtain a planar graph H that contains a subdivision of G and a subdivision of G^*. We can now embed H with straight-line segments in an $O(n) \times O(n)$ regular grid [41]. In this way we obtain an embedding of G and an embedding of G^* with the property that each edge and each dual edge is represented by a two-segment polygonal curve, and e and e^* cross as desired. With this embedding we can carry out actual operations using coordinates.

Vertices of G are usually denoted by x, y, u, v. Faces of G are usually denoted by symbols like f and g. The dual vertices are usually denoted using early letters of the Latin alphabet, like a and b. We use a_∞ for the dual vertex representing the outer face. We will denote cycles and paths in the dual graph with Greek letters, such as γ and π. Sets of cycles and paths in the dual graph are with capital Greek letters, like Γor Π.

Quite often we identify a graph object and its geometric representation in the embedding. In particular, (closed) walks in the graph define (closed) curves in the plane. We say that a closed walk γ in G^* is ***non-crossing*** if there is an infinitesimal perturbation γ_ε of the curve γ that makes it simple. If γ is simple, we can take $\gamma_\varepsilon = \gamma$. For each simple closed curve γ in the plane, let $\mathrm{int}(\gamma)$ be the bounded domain of $\mathbb{R}^2 \setminus \gamma$, and let $\mathrm{ext}(\gamma)$ be the unbounded one. For each closed, non-crossing closed walk γ in the dual graph G^*, let $V_{\mathrm{int}}(\gamma, G) = \mathrm{int}(\gamma_\varepsilon) \cap V(G)$ and $V_{\mathrm{ext}}(\gamma, G) = \mathrm{ext}(\gamma_\varepsilon) \cap V(G)$. Note that since γ is a walk in G^*, the vertices of $V(G)$ are far away from γ and it does not matter which infinitesimal perturbation γ_ε of γ we use. See Figure 2 for an example

Distances in graphs. In this paper we allow that the arcs have negative lengths λ. However, the graphs cannot have negative cycles, that is, cycles of negative length. In our approach we need that subpaths of shortest paths are also shortest paths.

For a graph G, a ***shortest-path tree from*** a vertex $r \in V(G)$ is a tree T that is a subgraph of G and satisfies $d_T(r, y) = d_G(r, y)$ for all $y \in V(G)$. A ***shortest-path tree to*** a vertex $r \in V(G)$ is a tree T that is a subgraph of G and satisfies $d_T(y, r) = d_G(y, r)$ for all $y \in V(G)$.

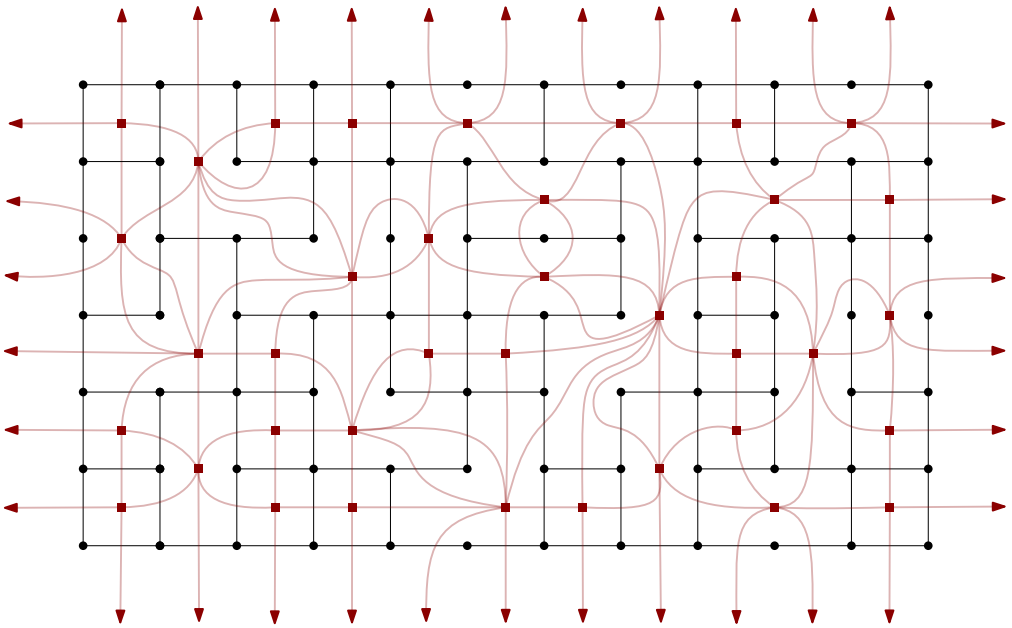

Figure 1: A plane graph G (in black with dots for the vertices) and its dual G^* (in red with squares for the vertices). The dual vertex a_∞ corresponding to the outer face of G is not drawn. Dual edges with an endpoint at a_∞ are represented using arrows.

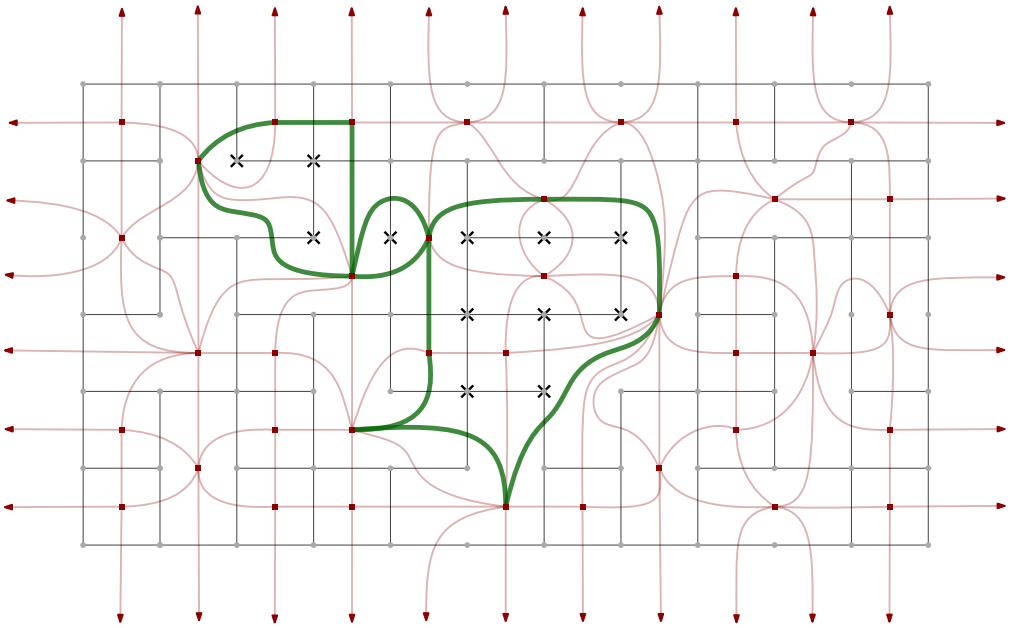

Figure 2: A non-crossing closed walk γ in the graph of Figure 1 is drawn in thick green. The vertices of $V_{\text{int}}(\gamma, G)$ are marked with crosses.

For all graphs considered in this paper we assume that, whenever we have an arc e, we also have its reversed arc e^R. We can ensure this by adding arcs with large enough length such that, if any shortest path uses them, this is detected. Similarly, adding edges, we can assume that the graphs that we are considering are connected.

For a given graph G with edge lengths $\lambda(\cdot)$, we use G^R for the reversed graph, that is, the graph G with edge lengths $\lambda^R(e) = \lambda(e^R)$. A shortest-path tree from r in G^R is the reversal of a shortest-path tree to r in G. Thus, as far as computation is concerned, there is no difference between computing shortest-path trees from or to a vertex.

Vertex-based information. Consider a graph G. For each vertex $x \in V(G)$, each subset $U \subseteq V(G)$, and each real value δ, we define

$$\begin{aligned} \operatorname{diam}(x,U,G) &:= \max\{d_G(x,u) \mid u \in U\}, \\ \operatorname{sum}(x,U,G) &:= \sum_{u\in U} d_G(x,u), \\ \operatorname{count}_{\leq}(x,U,G,\delta) &:= |\{u \in U \mid d_G(x,u) \leq \delta\}|. \end{aligned}$$

Our main results will compute these values for all vertices $x \in V(G)$ when G is planar and $U = V(G)$. Clearly we have

$$\begin{aligned} \operatorname{diam}(G) &= \max\{\operatorname{diam}(x,V(G),G) \mid x \in V(G)\}, \\ \operatorname{sum}(G) &= \sum_{x\in V(G)} \operatorname{sum}(x,V(G),G), \\ \operatorname{count}_{\leq}(G,\delta) &= \sum_{x\in V(G)} \operatorname{count}_{\leq}(x,V(G),G,\delta). \end{aligned}$$

3 Handling weights within a non-crossing walk

For the rest of this section, let G be a plane graph with n vertices. In this section we are not concerned with distances. Instead, we are concerned with vertex-weights. Assume that each vertex x of G has a weight $\omega(x) \in \mathbb{R}$. For each subset U of vertices and each value $\delta \in \mathbb{R}$ we define

$$\sigma(U) := \sum_{x\in U} \omega(x), \quad \mu(U) := \max_{x\in U} \omega(x), \quad \kappa_{\leq}(U,\delta) := |\{x \in U \mid \omega(x) \leq \delta\}|.$$

Let γ be a non-crossing closed walk in the dual graph G^*. We are interested in a way to compute $\sigma(V_{\text{int}}(\gamma,G))$ $\mu(V_{\text{int}}(\gamma,G))$, and $\kappa_{\leq}(V_{\text{int}}(\gamma,G),\delta)$ *locally*, after some preprocessing of G and G^*. Here, locally means that we would like to just look at the edges of γ. In the following, we assume that any non-crossing closed walk γ in G^* is traversed clockwise.

In the next section we concentrate on the computation of $\sigma(\cdot)$ and then explain how to use it to for computing $\kappa_{\leq}(\cdot,\delta)$. In Section 3.2 we discuss the computation of $\kappa_{\leq}(\cdot,\delta)$

3.1 Sum of weights and counting weights

We start adapting the approach by Park and Phillips [37] and Patel [38], which considered the computation of $\sigma(\cdot)$ when $\omega(x) = 1$ for all $x \in V(G)$. We summarize the ideas in the next lemma to make it self-contained. While most of the paper is simpler for undirected graphs, in the next lemma we do need the directed edges of the dual graph. We are not aware of a similar statement that would work using the undirected dual graph.

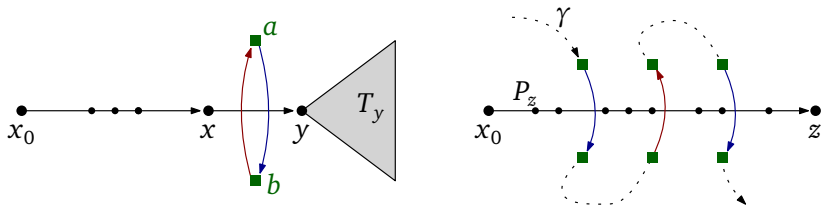

Figure 3: Proof of Lemma 1. Left: orientation of the edges of T. The dart $a \to b$ crosses $x \to y$ from left to right and thus $\chi(a \to b) = \omega(V(T_y))$. Right: the crossings of γ and P_z alternate between left-to-right and right-to-left, as we walk along P_z.

Lemma 1. *Let G be a plane graph, directed or not, and let x_0 be a fixed vertex in G. In linear time we can compute a weight function $\chi : E(G^*) \to \mathbb{R}$ with the following property. For every non-crossing closed walk γ in the dual graph G^* that is oriented clockwise and contains x_0 in its interior, we have*

$$\sigma(V_{\text{ext}}(\gamma, G)) = \sum_{a \to b \in \gamma} \chi(a \to b).$$

Proof. Take any spanning tree T of G rooted at x_0, and orient the arcs away from x_0. For example, a BFS tree of G from x_0. For each node $y \in V(G)$, let T_y be the subtree of T rooted at y. See Figure 3, left. For each vertex $y \neq x_0$ we proceed as follows. Let x be the parent of y and let $a \to b$ be the dual arc that crosses $x \to y$ from left to right. Then we assign $\chi(a \to b) = \sigma(V(T_y))$ and $\chi(b \to a) = -\chi(a \to b)$. For any dual edge ab of $E(G)^* \setminus E(T)^*$ we set $\chi(a \to b) = \chi(b \to a) = 0$. This finishes the description of the function χ. It is easy to see that we can compute χ in linear time.

From the definition of χ we have

$$\begin{aligned} \sum_{a \to b \in \gamma} \chi(a \to b) &= \sum_{a \to b \in E(\gamma) \cap E(T)^*} \chi(a \to b) \\ &= \sum_{\substack{x \to y \in T, \\ \gamma \text{ crosses } x \to y \\ \text{left-to-right}}} \sigma(V(T_y)) - \sum_{\substack{x \to y \in T, \\ \gamma \text{ crosses } x \to y \\ \text{right-to-left}}} \sigma(V(T_y)). \end{aligned} \tag{1}$$

Let γ_ε be an infinitesimal perturbation of γ that is simple. We then have $\text{int}(\gamma_\varepsilon) \cap V(G) = V_{\text{int}}(\gamma, G)$ and $\text{ext}(\gamma_\varepsilon) \cap V(G) = V_{\text{ext}}(\gamma, G)$.

Consider any vertex z of $V(G)$ and let P_z be the path in T from x_0 to z. Since x_0 is in $\text{int}(\gamma_\varepsilon)$ and γ_ε is a simple curve, the crossings between P_z and γ_ε, as we walk along along P_z, alternate between left-to-right and right-to-left crossings. See Figure 3, right. Since γ_ε defines a simple curve, the number of crossings is even if z is in $\text{int}(\gamma_\varepsilon)$ and odd otherwise. It follows that $\omega(z)$ contributes to the sum on the right side of equation (1) either once, if z is in $\text{ext}(\gamma_\varepsilon)$, or zero times, if z is in $\text{int}(\gamma_\varepsilon)$. The result follows. □

Lemma 1 can also be used to compute $\sigma(V_{\text{int}}(\gamma, G))$ because $\sigma(V_{\text{int}}(\gamma, G)) + \sigma(V_{\text{ext}}(\gamma, G) = \sigma(V(G))$.

We would like a data structure to quickly handle non-crossing closed walks in the dual graph that will be described compactly. More precisely, at preprocessing time we are given a family $\Pi = \{\pi_1, \ldots, \pi_\ell\}$ of walks in G^*, and the non-crossing closed walk will be given as a concatenation of some subwalks from Π. Using the function $\chi(\cdot)$ and partial sums over the edges e of each prefix of a walk in Π we get the following result.

Theorem 2. *Let G be a plane graph with n vertices and vertex-weights $\omega(\cdot)$. Let x_0 be a vertex of G. Let $\Pi = \{\pi_1, \ldots, \pi_\ell\}$ be a family of walks in G^* with a total of m edges, counted with multiplicity. After $O(n+m)$ preprocessing time, we can answer the following type of queries: given a non-crossing*

closed walk γ in G^*, described as a concatenation of k subpaths from Π, and with the property that γ is oriented clockwise and contains x_0 in its interior, return $\sigma(V_{\text{int}}(\gamma, G))$ in $O(k)$ time.

Proof. We compute for G the function χ of Lemma 1. For each walk π_i of Π we proceed as follows. Let $e(i,1),\ldots,e(i,m_i)$ be the arcs of π_i, as they appear along the walk π_i, and define the partial sums $S[i,j] = \sum_{t=1}^{j} \chi(e(i,t))$ for $j = 1,\ldots,m_i$. It is also convenient to define $S[i,0] = 0$. The values $S[i,1],\ldots,S[i,m_i]$ can be computed in $O(m_i)$ time using that $S[i,j] = S[i,j-1] + \chi(e(i,j))$ for $j = 1,\ldots,m_i$. Repeating the procedure for each $\pi_i \in \Pi$, we have spent a total of $O(n+m)$ time. This finishes the preprocessing.

Consider a non-crossing closed walk γ in G^* given as the concatenation of k walks $\pi^1,\ldots,\pi^k$, each of them a subpath of some path in Π. Each π^t in the description of γ is of the form $e(i(t),j_1(t)),\ldots,e(i(t),j_2(t))$ for some index $i(t)$ (so π^t is a subpath of $\pi_{i(t)}$) and some indices $j_1(t), j_2(t)$ with $1 \le j_1(t) \le j_2(t) \le m_{i(t)}$. Then we have

$$\sum_{e \in \pi^t} \chi(e) \;=\; S[i(t), j_2(t)] - S[i(t), j_1(t)-1].$$

Because of the properties of χ in Lemma 1 we have

$$\begin{aligned}\sigma(V_{\text{ext}}(\gamma, G)) \;&=\; \sum_{a\to b \in \gamma} \chi(a\to b) \;=\; \sum_{t=1}^{k} \sum_{a\to b \in \pi^t} \chi(a\to b) \\ &=\; \sum_{i=t}^{k} S[i(t), j_2(t)] - S[i(t), j_1(t)-1].\end{aligned}$$

It follows that we can compute $\sigma(V_{\text{ext}}(\gamma, G))$ in $O(k)$ time, and therefore we can also obtain $\sigma(V_{\text{int}}(\gamma, G))$ in the same time bound. □

We now look into the case of computing $\kappa_{\le}(\cdot, \delta)$. Using a binary search on $W = \{\omega(v) \mid v \in V(G)\}$ we achieve the following result. Note that in the following result the dependency in n increases.

Corollary 3. *Consider the setting of Theorem 2. After $O(n(n+m))$ preprocessing time, we can answer the following type of queries: given a value $\delta \in \mathbb{R}$ and a non-crossing closed walk γ in G^*, described as a concatenation of k subpaths from Π, and with the property that γ is oriented clockwise and contains x_0 in its interior, return $\kappa_{\le}(V_{\text{int}}(\gamma, G), \delta)$ in $O(k + \log n)$ time.*

Proof. We sort the n weights $W = \{\omega(v) \mid v \in V(G)\}$ and store them in an array. Let $w_1,\ldots,w_n$ be the resulting weights, so that $w_1 \le \cdots \le w_n$. For $i = 1,\ldots,n$, we define the weight function ω_i by

$$\omega_i(v) \;=\; \begin{cases} 1 & \text{if } \omega(v) \le w_i, \\ 0 & \text{otherwise.} \end{cases}$$

Then, we apply Theorem 2 for each of the weight functions $\omega_1,\ldots,\omega_n$. This finishes the preprocessing.

To compute $\kappa_{\le}(V_{\text{int}}(\gamma, G), \delta)$ for a given $\delta \in \mathbb{R}$, we make a a binary search in W to find $w_i = \max\{w \in W \mid w \le \delta\}$ and then use the data structure for the weight function ω_i to get

$$\begin{aligned}\sum_{v \in V_{\text{int}}(\gamma, G)} \omega_i(v) \;&=\; |\{v \in V_{\text{int}}(\gamma, G) \mid \omega(v) \le w_i\}| \\ &=\; \kappa_{\le}(V_{\text{int}}(\gamma, G), w_i) \\ &=\; \kappa_{\le}(V_{\text{int}}(\gamma, G), \delta).\end{aligned}$$

Thus, a query boils down to a (standard) binary search followed by a single query to the data structure of Theorem 2. Therefore the query time is $O(k + \log n)$. □

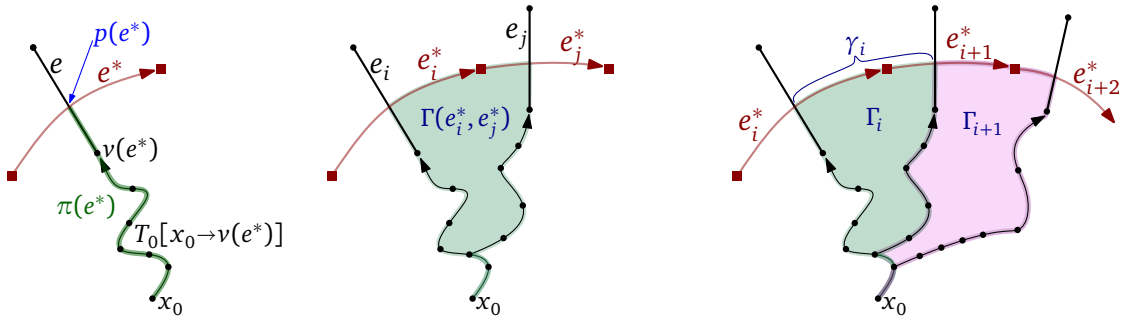

Figure 4: Defining $\chi_\mu(\cdot)$ in the proof of Lemma 4. Left: Notation for a dual edge e^*. The path $\pi(e^*)$ is thicker. Middle: The region $\Gamma(e_i^*, e_j^*)$. We mark the portion of T_0 that is included in $\Gamma(e_i^*, e_j^*)$. Right: the regions Γ_i.

3.2 Maximum weight

The proof of Lemma 1 heavily uses that the sum has an inverse operation. We are not aware of any such result for computing the maximum weight, $\mu(V_{\text{int}}(\gamma, G))$ or $\mu(V_{\text{ext}}(\gamma, G))$. We could do something similar as we did in the proof of Corollary 3, namely, a binary search in $W = \{\omega(v) \mid v \in V(G)\}$ to find the largest weight inside $V_{\text{int}}(\gamma, G)$. However, the extra preprocessing time in Corollary 3, as compared to the preprocessing time of Theorem 2, leads to a worst running time in our target application. We will develop now a different approach that works for a special type of closed walks that we have in our application.

Let x_0 be a vertex of G and let T_0 be a spanning tree of G rooted at x_0. We say that a cycle γ in the dual graph G^* is T_0-***star-shaped*** if the root x_0 is in $\text{int}(\gamma)$ and, for each vertex y in $V_{\text{int}}(\gamma, G)$, the whole path in T_0 from x_0 to y is contained in $\text{int}(\gamma)$. (Note that the concept is not meaningful for closed walks that repeat some vertex; hence our restriction to cycles for the time being.) We define the following family of dual cycles:

$$\Xi(G, T_0) \;=\; \{\gamma \mid \gamma \text{ is a } T_0\text{-star-shaped cycle in } G^*\}. \tag{2}$$

Lemma 4. *There is a weight function $\chi_\mu : E(G^*) \times E(G^*) \to \mathbb{R}$ with the following properties:*

- *For every cycle $\gamma = e_0^* e_1^*, \ldots, e_{k-1}^*$ of $\Xi(G, T_0)$ that is oriented clockwise*

$$\mu(V_{\text{int}}(\gamma, G)) \;=\; \max\left\{\chi_\mu(e_i^*, e_{i+1}^*) \mid i = 0, \ldots, k-1\right\} \qquad \textit{(indices modulo } k\textit{)}$$

- *A representation of the function χ_μ can be computed in linear time.*
- *For any two dual edges ab and bc of G^*, the value $\chi_\mu(ab, bc)$ can be obtained in constant time.*

Proof. In this proof, for each vertex v, we use $T_0[x_0 \to v]$ to denote the path in T_0 from x_0 to v.

For a dual arc e^*, let $p(e^*)$ be intersection point of e and e^*, let $v(e^*)$ be the vertex of e to the right of e^*, and let $\pi(e^*)$ be the curve obtained by the concatenation of $T_0[x_0 \to v(e^*)]$ and the portion of e from $v(e^*)$ to $p(e^*)$. See Figure 4, left.

We can now provide a definition of the function χ_μ. Consider any two dual edges e_i^* and e_j^* in G^*. If they have no common vertex or if they are equal, then we set $\chi_\mu(e_i^*, e_j^*) = 0$. This is not very relevant because such terms never show up in the desired properties. It remains to consider the case when they have a common vertex. For this case we then define $\Gamma(e_i^*, e_j^*)$ as the region of the plane bounded by $\pi(e_i^*)$, the portion of $e_i^* e_j^*$ from $p(e_i^*)$ to $p(e_j^*)$, and the reverse of $\pi(e_j^*)$. We regard the region $\Gamma(e_i^*, e_j^*)$ as a closed set, with its boundary and the curves that define it. See

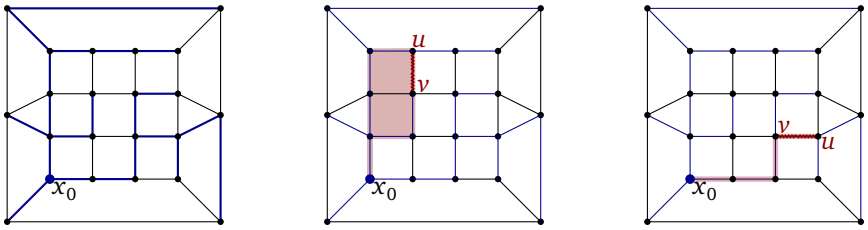

Figure 5: Example of the regions $R(uv)$. Left: a graph G and a spanning tree T_0. Center and right: the regions $R(uv)$ for two different edges uv.

Figure 4, center. Finally, we set $\chi_\mu(e_i^*, e_j^*) = \mu\Big(V(G) \cap \Gamma(e_i^*, e_j^*)\Big)$. We will discuss the efficient computation and representation of $\chi_\mu(e_i^*, e_j^*)$ later.

We claim that χ_μ satisfies the property in the first item. Consider any dual cycle $\gamma = e_0^* e_1^*, \ldots, e_{k-1}^*$ and let A be the closure of $\mathrm{int}(\gamma)$. For each i, let γ_i be the curve described by γ from $p(e_i^*)$ to $p(e_{i+1}^*)$ and let us use the shorthand $\Gamma_i = \Gamma(e_i^*, e_{i+1}^*)$, where indices are modulo k. See Figure 4, right. Note that γ_i is one of the curves used to define the region Γ_i. If γ is in $\Xi(G, T_0)$, then the three curves that bound Γ_i are contained in A and therefore Γ_i is contained in A $(i = 0, \ldots, k-1)$. Moreover, since $\gamma_0, \ldots, \gamma_{k-1}$ is a decomposition of γ and the regions Γ_i and Γ_{i+1} $(i = 0, \ldots, k-1)$ share the path $\pi(e_{i+1}^*)$ on its boundary, the union $\cup_i \Gamma_i$ is precisely A. (It seems that the union is interior-disjoint, but we do not need that.) Since the boundary of A does not contain any vertex of G, we get

$$\begin{aligned}
\mu(V_{\mathrm{int}}(\gamma, G)) &= \max\{\omega(v) \mid v \in V(G) \cap \mathrm{int}(\gamma)\} \\
&= \max\{\mu(V(G) \cap \Gamma_i) \mid i = 0, \ldots, k-1\} \\
&= \max\left\{\mu(V(G) \cap \Gamma(e_i^*, e_{i+1}^*) \mid i = 0, \ldots, k-1\right\} && \text{(indices modulo } k\text{)} \\
&= \max\left\{\chi_\mu(e_i^*, e_{i+1}^*) \mid i = 0, \ldots, k-1\right\}. && \text{(indices modulo } k\text{)}
\end{aligned}$$

We have shown that the property in the first item holds.

It remains to discuss the computational part. First we discuss an alternative definition of χ_μ that is more convenient for the computation. For each edge uv of G, let $R(uv)$ be the region of the plane defined by the paths in T_0 from x_0 to both endpoints of uv and the edge uv itself. We include in $R(uv)$ the two paths used to define it and the edge uv. If the paths in T_0 from x_0 to u and v share a part, then the region $R(uv)$ also contains that common part. If uv is in T_0, then the region $R(uv)$ is actually a path contained in T_0. See Figure 5 for an example. Finally, for each edge uv of G we define $\varphi(uv)$ as

$$\varphi(uv) := \mu(V(G) \cap R(uv)) = \max\{\omega(x) \mid x \in V(G) \cap R(uv)\}.$$

For each vertex v of G, we define $\varphi(v)$ as the maximum weight on the path $T_0[x_0 \to v]$. This last case can be interpreted as a degenerate case of the previous one. Indeed, if v' is the parent of v in T_0, then $\varphi(v) = \varphi(vv')$.

Let f be a face of G and let e_i and e_j be two edges on the boundary of f. We are going to give an alternative definition of $\chi_\mu(e_i^*, e_j^*)$. If e_j is *not* the follower of e_i along the counterclockwise traversal of f, let $E(f, e_i, e_j)$ be the edges between e_i and e_j in a counterclockwise traversal of f. We do not include e_i and e_j in $E(f, e_i, e_j)$, but the set $E(f, e_i, e_j)$ is nonempty by assumption. See Figure 6 for an illustration. In this case we have

$$\chi_\mu(e_i^*, e_j^*) = \max\{\varphi(e) \mid e \in E(f, e_i, e_j)\}. \tag{3}$$

To see that these equality indeed holds, note that the difference between the region $\Gamma(e_i^*, e_j^*)$ and $\bigcup\{R(e) \mid e \in E(f, e_i, e_j)\}$ is just a portion of the interior of the face f, which cannot contain vertices

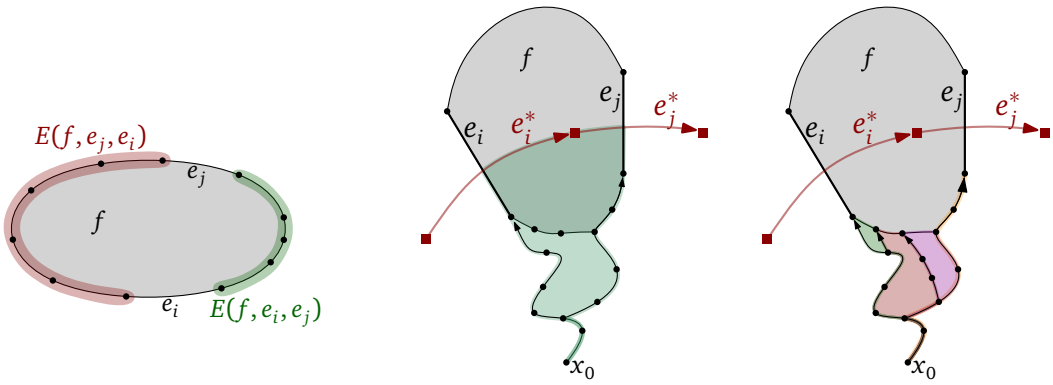

Figure 6: Left: example for the sets $E(f,e_i,e_j)$. Center and right: example indicating the relation between $\Gamma(e_i^*,e_j^*)$ and $\bigcup\{R(e) \mid e \in E(f,e_i,e_j)\}$. (Compare to Figure 4.) Note that on the right we have five regions $R(e)$, but two of them degenerate to paths on T_0.

of G. See Figure 6, center and right, for an illustration. If e_i and e_j are consecutive along the counterclockwise traversal of f, then they have a common vertex v and we have $\chi_\mu(e_i^*,e_j^*) = \varphi(v)$. The argument in this case is the same: the difference between $\Gamma(e_i^*,e_j^*)$ and the path $T_0[x_0 \to v]$ is a portion of the interior of f.

The second, alternative definition of χ_μ is more suitable for efficient management. First, we compute the values $\varphi(\cdot)$. For this we use the undirected version of G. Let $C = E(G) \setminus E(T_0)$ be the primal edges not contained in T_0. The dual of those edges, C^*, form a spanning tree of the dual graph G^*. The pair (T_0, C) is a so-called tree-cotree decomposition. We root C^* at the dual vertex representing the outer face of G. Each edge $e \in C$ defines a region A_e of the plane, namely the closed region bounded by the unique simple closed curve contained in $T_0 + e$. Note that, for each $uv \in C$, the region $R(uv)$ is precisely the union of A_{uv} and the two paths in T_0 from x_0 to the endpoints of uv. Each edge $e \in C$ defines a dual subtree, denoted by C_e^*, which is the component of $C^* - e^*$ without the root. The region A_e corresponds to the faces of G that dualize to vertices of C_e^*. See Figure 7 for an example. After computing and storing for each face of G the maximum weight of its incident vertices, we can use a bottom-up traversal of the dual tree C^* and the values stored for each face to compute the values $\mu(A_e \cap V(G))$ in linear time for all edges $e \in C$. With a top-bottom traversal of the primal tree T_0 we can also compute and store for each node v of G the values $\mu(T_0[x_0 \to v] \cap V(G))$. From this we can compute $\varphi()$ as follows:

$$\begin{aligned}
\forall v \in V(G): &\quad \varphi(v) = \mu(T_0[x_0 \to v] \cap V(G)),\\
\forall uv \in E(T_0): &\quad \varphi(uv) = \max\{\mu(T_0[x_0 \to v] \cap V(G)), \omega(u), \omega(v)\},\\
\forall uv \in C: &\quad \varphi(uv) = \max\{\mu(A_{uv} \cap V(G)), \mu(T_0[x_0 \to u] \cap V(G)), \mu(T_0[x_0 \to v] \cap V(G))\}.
\end{aligned}$$

Since each value on the right side is already computed, we spend linear time to compute the values $\varphi(\cdot)$.

To represent χ_μ compactly, we will use a data structure for range minimum queries: preprocess an array of numbers $A[1 \ldots m]$ such that, at query time, we can report $\min\{A[k] \mid i \le k \le j\}$ for any given query pair of indices $i < j$. There are data structures that use linear-time preprocessing and $O(1)$ time per query; see [2, 15] for a description. This data structure does exploit the full power of random-access memory (RAM). It is trivial to extend this data structure for circular arrays: each query in a circular array corresponds to two queries in a linear array.

For each face f of G, we build a circular array $A_f[\cdot]$ indexed by the edges, as they appear along the face f. At the entry $A_f[e]$ we store the value $\varphi(e)$. For each face we spend time proportional to the number of edges on the boundary of the face. Thus, for the whole graph G this preprocessing takes linear time. For two edges e_1 and e_2 on the boundary of a face f and

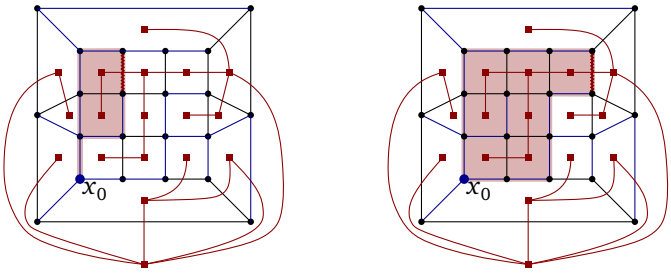

Figure 7: Dual tree defined by the cotree C and its relevance to compute $\varphi(uv)$. Here we show the relevance for two different edges $uv \in (E(G) \setminus E(T_0))^*$ in the example of Figure 7.

with no common vertex, the value $\chi_\mu(e_1^*, e_2^*)$, as described in (3), is precisely a range maximum query in the circular array $A_f[\cdot]$, and thus can be answered in constant time. The case when e_1 and e_2 have a common vertex is easier because for each edge e_1 there are only two such possible edges e_2, one per face with e_1 is on the boundary. □

For our application we will have to deal with pieces that have holes and thus a part of T_0 may be missing. Because of this, we also need to extend things to a type of non-crossing walks.

Like before, let G be a plane graph and let T_0 be a rooted spanning subtree. Let P be a subgraph of G, with the embedding inherited from G. Assume that the root x_0 of T_0 is in P. A non-crossing closed walk γ in P^* is T_0-***star-shaped*** if the root x_0 is in $\mathrm{int}(\gamma)$ and, for each vertex y in $V_{\mathrm{int}}(\gamma, P)$, all the vertices of P in the path $T_0[x_0 \to y]$ are contained in $\mathrm{int}(\gamma)$. We can define the following family of dual non-crossing walks:

$$\widetilde{\Xi}(G, P, T_0) \;=\; \{\gamma \mid \gamma \text{ is a } T_0\text{-star-shaped non-crossing walk in } P^*\}. \tag{4}$$

Theorem 5. *Let G be a plane graph with n vertices and vertex-weights $\omega(\cdot)$, and let T_0 be a rooted spanning tree in G. Let P be a subgraph of G vertices that contains the root of T_0. Let b be the number of edges on all the faces of P that contain edges of $E(T_0) \setminus E(P_0)$. Let $\Pi = \{\pi_1, \ldots, \pi_\ell\}$ be a family of walks in P^* with a total of m edges, counted with multiplicity. After $O(n + m + b^3)$ preprocessing time, we can answer the following type of queries: given a closed walk γ in $\widetilde{\Xi}(G, P, T_0)$, described as a concatenation of k subpaths from Π and oriented clockwise return $\mu(V_{\mathrm{int}}(\gamma, P))$ in $O(k)$ time.*

Although the dependency on b in the time bound can perhaps be reduced, it is enough for our purposes because currently the bottleneck is somewhere else.

Proof. We may assume that $V(G) = V(P) = V(T_0)$. To see this, first we note that we can remove edges of G that are not in $E(T_0) \cup E(P)$ because they do not play any role. Then we can replace each maximal subtree of $T_0 - V(P)$ by edges that connect vertices of P without changing the set $\widetilde{\Xi}(G, P, T_0)$. For this we just need the ancestor-descendant relation between vertices incident to the face. See Figure 8 to see the transformation. Thus, from now on we restrict ourselves to the case where $V(G) = V(P) = V(T_0)$.

Let F be the set of faces of P that contain some edges of $E(T_0) \setminus E(P)$ and consider the set

$$A \;=\; \{(e_1, e_2) \in E(P)^2 \mid e_1, e_2 \in E(f) \text{ for some } f \in F\}.$$

It is clear that A has $O(b^2)$ pairs. For each $(e_1, e_2) \in A$, let f be the face of F that have e_1 and e_2 on the boundary and compute a dual path $\pi_G(e_1^*, e_2^*)$ in G^* from e_1^* to e_2^* whose other edges are contained in the face f. This means that $\pi_G(e_1^*, e_2^*)$, besides e_1^* and e_2^*, only uses edges of $E(T_0)^* \setminus E(P)^*$. (If e_1 and e_2 are cofacial in G, then $\pi_G(e_1^*, e_2^*) = e_1^* e_2^*$.) See Figure 9 for an example. In particular, since $E(T_0) \setminus E(P)$ is a forest on b vertices, it has at most b edges, and the

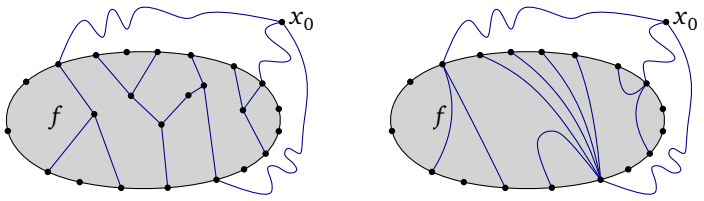

Figure 8: Transformation in the proof of Theorem 5 to assume that $V(G) = V(P) = V(T_0)$. The vertices of $V(T_0)$ inside a face f of P^* can be removed and we use direct edges representing the paths in T_0.

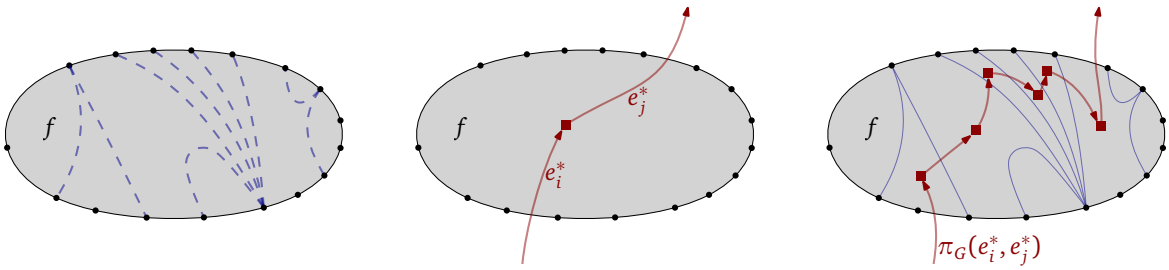

Figure 9: Construction of the paths $\pi_G(e_1^*, e_2^*)$. Left: a face of P with edges of $T_0 - E(P)$ dashed. Center: a two-edge walk in P^*. Right: the corresponding path $\pi_G(e_1^*, e_2^*)$ in G^*..

path $\pi_G(e_1^*, e_2^*)$ has $O(b)$ edges. Thus, the paths $\pi_G(e_1^*, e_2^*)$, over all elements $(e_1, e_2) \in A$, have together $O(b^3)$ edges.

The paths $\{\pi_G(e_1^*, e_2^*) \mid (e_1, e_2) \in A\}$ are used to naturally transform walks in P^* into walks in G^*. Indeed, if we have a walk α in P^* and we replace each occurrence of $e_1^* e_2^*$, where $(e_1, e_2) \in A$ by $\pi_G(e_1^*, e_2^*)$, then we obtain a walk in G^*.

We compute for G the function χ_μ of Lemma 4. For each element $(e_1, e_2) \in A$, we compute and store

$$\varphi(e_1^*, e_2^*) \;=\; \max\{\chi_\mu(ab, bc) \mid ab \text{ and } bc \text{ consecutive dual edges along } \pi_G(e_1^*, e_2^*)\}.$$

Using the properties of χ_μ stated in Lemma 4 and using that the paths $\{\pi_G(e_1^*, e_2^*) \mid (e_1, e_2) \in A\}$ have $O(b^3)$ edges, we can do this step in $O(n + b^3)$ time.

For each walk π_i of Π we proceed as follows. Let $e_1^*, \ldots, e_{m_i}^*$ be the edges of π_i, as they appear along π_i. We make an array $A_i[1..(m_i - 1)]$ such that

$$A_i[j] \;=\; \begin{cases} \varphi(e_j^*, e_{j+1}^*) & \text{if } (e_j, e_{j+1}) \in A, \\ \chi_\mu(e_j^*, e_{j+1}^*) & \text{if } (e_j, e_{j+1}) \notin A. \end{cases}$$

Finally, we store the array $A_i[\cdot]$ for range maximum queries [2]; see the discussion at the end of the proof of Lemma 4. We spend $O(m_i)$ preprocessing time for π_i and can find $\min A[j..j']$ in constant time for any given indices $1 \le j < j' < m_i$. This step, together for all paths $\pi_i \in \Pi$, takes $O(\sum_i m_i) = O(m)$ time. This finishes the preprocessing.

Assume that we are given a non-crossing closed walk γ in $\widetilde{\Xi}(G, P, T_0)$, given as the concatenation of k paths $\pi^1, \ldots, \pi^k$, each of them a subpath of some path in Π. Let γ_G be the closed walk obtained from γ as follows: for each $(e_1, e_2) \in A$ and each appearance of $e_1^* e_2^*$ in γ, we replace $e_1^* e_2^*$ by $\pi_G(e_1^*, e_2^*)$. Note that γ_G is closed walk in G^*. In fact, γ_G it is a cycle in G^* because geometrically each single replacement occurs within a single face of F and all the replacements within a face do not introduce crossings because γ was non-crossing. Moreover, because each replacement is a rerouting within a face F and $V(G) = V(P)$, we have $V_{\text{int}}(\gamma, P) = V_{\text{int}}(\gamma_G, G)$. From Lemma 4 we

thus get that

$$\begin{aligned} V_{\text{int}}(\gamma, P) &= \mu(V_{\text{int}}(\gamma, G)) \\ &= \max\{\chi_\mu(ab, bc) \mid ab \text{ and } bc \text{ consecutive dual edges along the cycle } \gamma_G\}. \end{aligned}$$

Like in the proof of 2, we can break the computation of $\chi_\mu(\cdot)$ for pairs of consecutive edges of γ_G into k parts that occur within some path $\pi_i \in \Pi$ (after replacements) and k parts that use the last edge of π^t and the first of π^{t+1} ($t = 0, \ldots, k$, indices modulo k). The part within a path $\pi_i \in \Pi$ can be retrieved in constant time from the range maximum query for $A_i[\cdot]$. The part combining consecutive subpaths can be computed in constant time, but we have two cases to consider. Let ab be the last dual edge of π^t and let bc be the first dual edge of π^{t+1}. If $(ab, bc) \in A$, then we have to use $\varphi(ab, bc)$. Otherwise we can directly use $\chi_\mu(ab, bc)$, which can be computed in constant time (third item of Lemma 4). Finally, we have to take the maximum from those $2k$ values. □

When $G = P$, then $\widetilde{\Xi}(G, P, T_0) = \Xi(G, T_0)$, we have $b = 0$, and Theorem 5 simplifies to the following.

Corollary 6. *Let G be a plane graph with n vertices and vertex-weights $\omega(\cdot)$, and let T_0 be a rooted spanning tree in G. Let $\Pi = \{\pi_1, \ldots, \pi_\ell\}$ be a family of paths in G^* with a total of m edges, counted with multiplicity. After $O(n+m)$ preprocessing time, we can answer the following type of queries: given a cycle γ in $\Xi(G, T_0)$, described as a concatenation of k subpaths from Π and oriented clockwise return $\mu(V_{\text{int}}(\gamma, G))$ in $O(k)$ time.*

4 Abstract Voronoi diagrams

Abstract Voronoi diagrams were introduced by Klein [27] as a way to handle together several of the different types of Voronoi diagrams that were appearing. The concept is restricted to the plane $\mathbb{R}^2$. They are defined using the concept of *bisectors* and *dominant regions*. We will use the definition by Klein, Langetepe and Nilforoushan [29], as it seems the most recent and general. For the construction, we use the randomized incremental construction of Klein, Mehlhorn and Meiser [30], also discussed in [29] for their framework. In our notation, we will introduce an A in front to indicate we are talking about objects in the *abstract* Voronoi diagram.

Let S be a finite set, which we refer to as *abstract sites*. For each ordered $(p, q) \in S^2$ of distinct sites, we have a simple planar curve $\text{AJ}(p, q)$ and an open domain $\text{AD}(p, q)$ whose boundary is $\text{AJ}(p, q)$. We refer to the pair $(\text{AJ}(p, q), \text{AD}(p, q))$ as an ***abstract bisector***. Define for each $p \in S$ the ***abstract Voronoi region*** $\text{AVR}(p, S) = \bigcap_{q \in S \setminus \{p\}} \text{AD}(p, q)$. Then the ***abstract Voronoi diagram*** of S, denoted by $\text{AVD}(S)$, is defined as $\text{AVD}(S) = \mathbb{R}^2 \setminus \bigcup_{p \in S} \text{AVR}(p, S)$.

The intuition is that the set $\text{AD}(p, q)$ is the set of points that are closer to p than to q and that $\text{AJ}(p, q)$ plays the role of bisector. Then, $\text{AVR}(p, S)$ stands for the points that are dominated by p, when compared against all $q \in S \setminus \{p\}$. Note that $\text{AVR}(p, S)$ is an open set because it is the intersection of open sets. The abstract Voronoi diagram, $\text{AVD}(S)$ would then be the set of points where no site dominates, meaning that at least two sites are "equidistant" from the point. However, the theory does not rely on any such interpretations. This makes it very powerful but less intuitive: some arguments become more cumbersome.

While these concepts can be considered in all generality, the theory is developed assuming that certain properties, called axioms, are satisfied. A system of abstract bisectors $\{(\text{AJ}(p, q), \text{AD}(p, q)) \mid p, q \in S, p \neq q\}$ is ***admissible*** if it satisfies the following properties:

(A1) For all distinct $p, q \in S$, $J(p, q) = J(q, p)$.

(A2) For all distinct $p, q \in S$, the plane $\mathbb{R}^2$ is the disjoint union of $D(p, q)$, $J(p, q)$ and $D(q, p)$.

(A3) There exists a special point in the plane, which we call p_∞, such that, for all distinct $p, q \in S$, the curve $J(p,q)$ passes through p_∞.[1]

(A4) For each subset S' of S with 3 elements and each $p \in S'$, the abstract Voronoi region $\mathrm{AVR}(p,S')$ is path connected.

(A5) For each subset S' of S with 3 elements we have $\mathbb{R}^2 = \bigcup_{p\in S'} \overline{\mathrm{AVR}(p,S')}$.

For the rest of the discussion on abstract Voronoi diagrams, we assume that these axioms are satisfied. Note that axioms (A4)-(A5) are not the ones given in the definition of [29] but, as they show in their Theorem 15, they are equivalent. In this regard, our definition is closer to the one given in [28]. Since we are going to work with very natural, no-pathological Voronoi diagrams, any of the sets of axioms used in any of the other papers we have encountered also works in our case. Assuming these axioms, one can show that the abstract Voronoi diagram $\mathrm{AVD}(S)$ is a plane graph [29, Theorem 10]. This brings a natural concept of ***abstract Voronoi vertex*** and ***abstract Voronoi edge*** as those being vertices (of degree ≥ 3) and edges in the plane graph $\mathrm{AVD}(S)$.

Klein, Mehlhorn and Meiser provide a randomized incremental construction of abstract Voronoi diagrams. One has to be careful about what it means to compute an abstract Voronoi diagram, since it is not even clear how the input is specified. For their construction, they assume as primitive operation that one can compute the abstract Voronoi diagram of any five abstract sites. The output is combinatorially described with a plane graph H and telling the description of each vertex and edge of H. The description of a vertex or an edge is a pointer to a vertex or an edge, respectively, in the abstract Voronoi diagram for at most four abstract sites. Thus, we tell that an edge e of H corresponds to some precise abstract edge e' of $\mathrm{AVD}(S')$, where $|S'| \leq 4$. Whether $\mathrm{AVD}(S')$ can be computed explicitly or not, it depends on how the input bisectors can be manipulated.

Klein, Mehlhorn and Meiser consider a special case, which is the one we will be using, where the basic operation requires the abstract Voronoi diagram of only four sites. (This particular case is not discussed in [29], but they discuss the general case.)

Theorem 7 (Klein, Mehlhorn and Meiser [30])**.** *Assume that we have an admissible system of abstract bisectors for a set S of m sites. The abstract Voronoi diagram of S can be computed in $O(m \log m)$ expected time using an expected number of $O(m \log m)$ elementary operations. If the abstract Voronoi diagram of any three sites contains at most one abstract Voronoi vertex, besides the special point p_∞, then an elementary operation is the computation of an abstract Voronoi diagram for four sites.*

5 Voronoi diagrams in planar graphs

We will need additively weighted Voronoi diagrams in *plane* graphs. We first define Voronoi diagrams for arbitrary graphs. Then we discuss a representation using the dual graphs that works only for plane graphs and discuss some folklore properties. See for example the papers of Marx and Pilipczuk [32] or Colin de Verdière [10] for similar intuitions. The dual representation is the key to be able to use the machinery of abstract Voronoi diagrams as a black box.

5.1 Arbitrary graphs

Let G be an arbitrary graph, not necessarily planar, with no negative cycles. A ***site*** s is a pair (v_s, w_s), where $v_s \in V(G)$ is its ***location***, and $w_s \in \mathbb{R}$ is its ***weight***, possibly negative. With a slight

[1] Usually the axiom tells that the stereographic projection to the sphere of the curve $J(p,q)$ can be completed to a closed Jordan curve passing through the north pole. For us it will be more convenient to project from a different point and complete all curves within the plane to make them pass through p_∞.

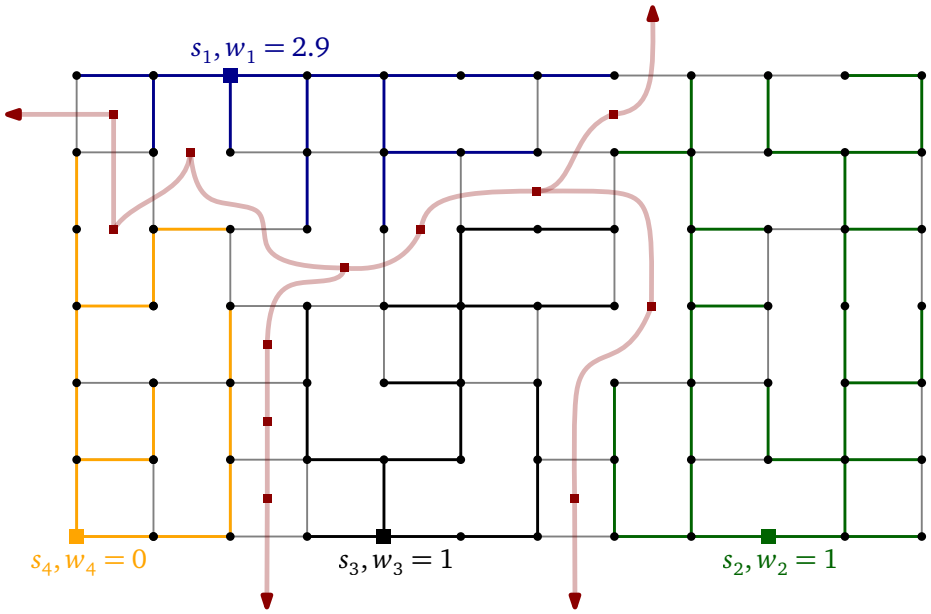

Figure 10: A graphic Voronoi diagram for four sites. The edges are undirected and have unit weight. An abstract Voronoi edge is marked with thicker pen.

abuse of notation, we will use s instead of v_s as the vertex. For example, for a site s we will write $s \in V(G)$ instead of $v_s \in V(G)$ and $d_G(s,x)$ instead $d_G(v_s,x)$.

Let S be a set of sites in G. For each $s \in S$, its ***graphic Voronoi region***, denoted $\mathrm{GVR}_G(s,S)$, is defined by

$$\mathrm{GVR}_G(s,S) \;=\; \{x \in V(G) \mid \forall t \in S \setminus \{s\} :\; w_s + d_G(s,x) \le w_t + d_G(t,x)\}.$$

See Figure 10 for an example. Note that we are using the distance from the sites to the vertices to define the graphic Voronoi cells. For directed graphs, using the reverse distance from the vertex to the sites would define different graphic regions (in general). However, this is equivalent to use the reversed graph G^R of G.

Even assuming that all distances in G are distinct, we may have $w_s + d_G(s,x) = w_t + d_G(t,x)$ for some vertex x. Also, some Voronoi cells may be empty. In our case, we will only deal with cases where these two things cannot happen. We say that the set S of sites is ***generic*** when, for each $x \in V(G)$ and for each distinct $s,t \in S$, we have $w_s + d_G(s,x) \neq w_t + d_G(t,x)$. The set S is ***independent*** when each Voronoi cell is nonempty. It is easy to see that, if S is a generic, independent set of sites, then $s \in \mathrm{GVR}_G(s,S)$ and each vertex x of $V(G)$ belongs to precisely one graphic Voronoi cell $\mathrm{GVR}_G(s,S)$ over all $s \in S$.

The ***graphic Voronoi diagram*** of S (in G) is the collection of graphic Voronoi regions:

$$\mathrm{GVD}_G(S) \;=\; \{\mathrm{GVR}_G(s,S) \mid s \in S\}.$$

The following property is standard.

Lemma 8. *Let S be a generic, independent set of sites. Then for each $s \in S$ it holds the following:*

- *For each x in $\mathrm{GVR}_G(s,S)$, the shortest path from s to x is contained in $\mathrm{GVR}_G(s,S)$.*
- *$\mathrm{GVR}_G(s,S)$ induces a connected subgraph of G.*

Proof. Let x be a vertex of $\mathrm{GVR}_G(s,S)$ and let $P(s,x)$ be the shortest path in G from s to x. Assume, for the sake of reaching a contradiction, that some vertex y on $P(s,x)$ is contained in some other

Voronoi cell $\mathrm{GVR}_G(t,S)$, where $t \neq s$. Because of uniqueness of shortest paths, this means that $d_G(t,y) < d_G(s,y)$. However, this implies that

$$d_G(t,x) \leq d_G(t,y) + d_G(y,x) < d_G(s,y) + d_G(y,x) = d_G(s,x),$$

where in the last equality we have used that y lies in the shortest path $P(s,x)$. The obtained inequality $d_G(t,x) < d_G(s,x)$ contradicts the property that $x \in \mathrm{GVR}_G(s,S)$. This proves the first item.

To show the second item, note that the subgraph of G induced $\mathrm{GVR}_G(s,S)$ contains (shortest) paths from s to all vertices of $\mathrm{GVR}_G(s,S)$ because of the previous item. □

For each two sites s and t, we define the ***graphic dominance region*** of s over t as

$$\begin{aligned}\mathrm{GD}_G(s,t) &= \mathrm{GVR}_G(s,\{s,t\}) \\ &= \{x \in V(G) \mid w_s + d_G(s,x) \leq w_t + d_G(t,x)\}.\end{aligned}$$

Lemma 9. *For each $s \in S$ we have $\mathrm{GVR}_G(s,S) = \bigcap_{t \in S\setminus\{s\}} \mathrm{GD}_G(s,t)$.*

Proof. We note that

$$\begin{aligned}\mathrm{GVR}_G(s,S) &= \{x \in V(G) \mid \forall t \in S \setminus \{s\} : w_s + d_G(s,x) \leq w_t + d_G(t,x)\} \\ &= \bigcap_{t \in S\setminus\{s\}} \{x \in V(G) \mid w_s + d_G(s,x) \leq w_t + d_G(t,x)\} \\ &= \bigcap_{t \in S\setminus\{s\}} \mathrm{GD}_G(s,t).\end{aligned}$$

□

5.2 Plane graphs

Now we will make use of graph duality to provide an alternative description of additively weighted Voronoi diagrams in plane graphs. The aim is to define Voronoi diagrams geometrically using bisectors, where a bisector is just going to be a cycle in the dual graph.

Consider two sites s and t in G and define

$$E_G(s,t) = \{xy \in E(G) \mid x \in \mathrm{GD}_G(s,t),\ y \in \mathrm{GD}_G(t,s)\}.$$

Thus, we are taking the edges that have each endpoint in different graphic Voronoi regions of $\mathrm{GVD}_G(\{s,t\})$. We denote by $E^*_G(s,t)$ their dual edges.

Lemma 10. *Let $\{s,t\}$ be a generic and independent set of sites. Then the edges of $E^*_G(s,t)$ define a cycle γ in G^*. Moreover, if $s \in V_{\mathrm{int}}(\gamma,G)$, then $V_{\mathrm{int}}(\gamma,G) = \mathrm{GD}_G(s,t)$ and $V_{\mathrm{ext}}(\gamma,G) = \mathrm{GD}_G(t,s)$.*

Proof. Let A^* be an arbitrary set of dual edges. It is well known that A^* is the edge set of a cycle if and only if $G - A$ has precisely two connected components. Moreover, two faces u^* and v^* of G^* are in the same side of the cycle defined by A^* if and only if u and v are in the same connected component of $G - A$. See for example the proof in [12, Proposition 4.6.1] or [4, Theorem 10.16].

When $\{s,t\}$ is generic and independent, we have $\mathrm{GD}_G(s,t) \neq \emptyset$, $\mathrm{GD}_G(s,t) \neq \emptyset$, and $V(G)$ is the disjoint union of $\mathrm{GD}_G(s,t)$ and $\mathrm{GD}_G(t,s)$, . This means that $E_G(s,t)$ is the edge cut between $\mathrm{GD}_G(s,t)$ and its complement, $\mathrm{GD}_G(t,s)$. Moreover, by Lemma 8, the subgraphs of G induced by $\mathrm{GD}_G(s,t)$ and by $\mathrm{GD}_G(t,s)$ are connected. Therefore $G - E_G(s,t)$ has precisely two connected components, and thus $E^*_G(s,t)$ is the edge set of a cycle γ in G^*.

Assume that $s \in V_{\mathrm{int}}(\gamma,G)$. Since $\mathrm{GD}_G(s,t)$ is the vertex set of the connected component of $G - E_G(s,t)$ that contains s, the faces of $\{u^* \mid u \in \mathrm{GD}_G(s,t)\}$ are in $\mathrm{int}(\gamma)$ and the faces $\{v^* \mid v \in \mathrm{GD}_G(t,s)\}$ are in $\mathrm{ext}(\gamma)$. Since a vertex u of G is the unique vertex of G contained in the dual face u^* of G^*, the result follows. □

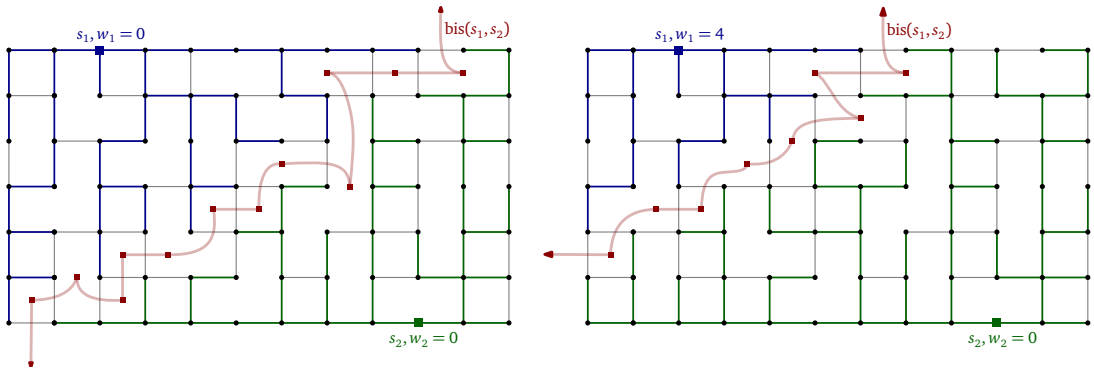

Figure 11: Two bisectors for sites placed at the same vertices but different weights.

When s and t are independent and generic, we define the ***bisector*** of s and t, denoted as $\text{bis}_G(s,t)$, as the curve in the plane defined by the cycle of $E^*_G(s,t)$, as guaranteed in the previous lemma. See figure 11 for an example. We also define $D_G(s,t)$ as the connected part of $\mathbb{R}^2 \setminus \text{bis}_G(s,t)$ that contains s. We then have

$$D_G(s,t) = \left(\bigcup_{v \in \text{GD}_G(s,t)} \overline{v^*} \right)^\circ . \tag{5}$$

Note that the pair $(\text{bis}_G(s,t), D_G(s,t))$ is the type of pair used to define abstract Voronoi diagrams. From now on, whenever we talk about the abstract Voronoi diagram of G, we refer to the abstract Voronoi diagram defined by the system of bisectors $\{(\text{bis}_G(s,t), D_G(s,t)) \mid s,t \in S, s \neq t\}$.

We have defined Voronoi regions of plane graphs in two different ways: using distances in the primal graph G, called *graphic* Voronoi regions, and using bisectors defined as curves in the plane, called *abstract* Voronoi regions. We next make sure that the definitions match, when restricted to vertices of G.

Lemma 11. *Let G be a plane graph and let S be a generic, independent set of sites. Then, for each $s \in S$, we have $\text{GVR}_G(s,S) = V(G) \cap \text{AVR}(s,S)$.*

Proof. Recall the definition

$$\text{AVR}(s,S) = \bigcap_{t \in S \setminus \{s\}} D_G(s,t).$$

Because of equation (5) we have

$$D_G(s,t) = \left(\bigcup_{v \in \text{GD}_G(s,t)} \overline{v^*} \right)^\circ ,$$

and we obtain that

$$\begin{aligned} \text{AVR}(s,S) &= \bigcap_{t \in S \setminus \{s\}} \left(\bigcup_{v \in \text{GD}_G(s,t)} \overline{v^*} \right)^\circ = \left(\bigcup_{v \in \bigcap_{t \in S \setminus \{s\}} \text{GD}_G(s,t)} \overline{v^*} \right)^\circ \\ &= \left(\bigcup_{v \in \text{GVR}_G(s,S)} \overline{v^*} \right)^\circ , \end{aligned}$$

where in the last equality we used Lemma 9. Since the only vertex of $V(G)$ contained in the dual face $\overline{v^*}$ is precisely v, and it lies in the interior of v^*, we get that $V(G) \cap \text{AVR}(s,S) = \text{GVR}(s,S)$. □

We cannot use the machinery of abstract Voronoi diagrams for arbitrary sites because of axiom (A3). In our case bisectors may not pass through a common "infinity point" p_∞. Indeed, for arbitrary planar graphs we could have two bisectors that never intersect. However, we can use it when all the sites are in the outer face of G. We next show this.

Lemma 12. *Let G be a plane graph and let S be a generic, independent set of sites located in the outer face of G. Let a_∞ be the vertex of G^* dual to the outer face of G. Then the system of abstract bisectors $\{(\mathrm{bis}_G(s,t), D_G(s,t)) \mid s,t \in S, s \neq t\}$ is admissible, where a_∞ plays the role of p_∞ in axiom (A3).*

Proof. It is clear that the system of abstract bisectors $\{(\mathrm{bis}_G(s,t), D_G(s,t)) \mid s,t \in S, s \neq t\}$ satisfies axioms (A1) and (A2) of the definition.

We next show the validity of axiom (A3). Consider any two sites s and t of S. Since $\mathrm{GVR}_G(s,S)$ and $\mathrm{GVR}_G(t,S)$ are nonempty, also $\mathrm{GD}_G(s,t)$ and $\mathrm{GD}_G(t,s)$ are nonempty. Since s and t are located in the outer face of G the bisector $\mathrm{bis}_G(s,t)$ passes through a_∞. Indeed, the dual faces s^* and t^* have to be in different sides of the dual cycle $\mathrm{bis}_G(s,t)$ and, since s and t are on the outer face of G, that can happen only if $\mathrm{bis}_G(s,t)$ passes through a_∞. Thus, if we take the geometric position of a_∞ as p_∞, all the bisecting curves pass through p_∞ and axiom (A3) holds.

For axiom (A4), consider any three sites r,s,t of S and let $S' = \{r,s,t\}$. Following the definition of $\mathrm{AVR}(\cdot,\cdot)$ we have

$$\mathrm{AVR}(s,S') = D_G(s,r) \cap D_G(s,t).$$

Because of equation (5), we have

$$\begin{aligned}
\mathrm{AVR}(s,S') &= \left(\bigcup_{v \in \mathrm{GD}_G(s,r)} \overline{v^*} \right)^\circ \bigcap \left(\bigcup_{v \in \mathrm{GD}_G(s,t)} \overline{v^*} \right)^\circ \\
&= \left(\bigcup_{v \in \mathrm{GD}_G(s,r) \cap \mathrm{GD}_G(s,t)} \overline{v^*} \right)^\circ \\
&= \left(\bigcup_{v \in \mathrm{GVR}_G(s,S')} \overline{v^*} \right)^\circ,
\end{aligned}$$

where in the last equality we have used Lemma 9. Since the vertices of $\mathrm{GVR}_G(s,S')$ form a connected subgraph of G (Lemma 8), the domains $\overline{v^*}$, when v iterates over $\mathrm{GVR}_G(s,S')$, are glued through the primal edges, and $\mathrm{AVR}(s,S')$ is path connected. This proves axiom (A4).

Axiom (A5) is shown similarly. Following the notation and the observations from the previous paragraph, we use that

$$\overline{\mathrm{AVR}(s,S')} = \overline{\bigcup_{v \in \mathrm{GVR}_G(s,S')} v^*}$$

and that $V(G) = \mathrm{GVR}_G(r,S') \cup \mathrm{GVR}_G(s,S') \cup \mathrm{GVR}_G(t,S')$, to conclude that

$$\overline{\mathrm{AVR}(r,S')} \cup \overline{\mathrm{AVR}(s,S')} \cup \overline{\mathrm{AVR}(t,S')} = \overline{\bigcup_{v \in V(G)} v^*} = \mathbb{R}^2.$$

□

The abstract Voronoi diagram $\mathrm{AVD}(S)$ is a plane graph, and by construction it is contained in the dual graph G^*. An abstract Voronoi vertex corresponds to a vertex in the dual graph G^*. An abstract Voronoi edge corresponds to a path in the dual graph G^*. More precisely, any abstract Voronoi edge corresponds to a portion of a bisector $\mathrm{bis}_G(s,t)$ whose endpoints are vertices of G^*.

We further have the following observation regarding the structure of abstract Voronoi diagrams.

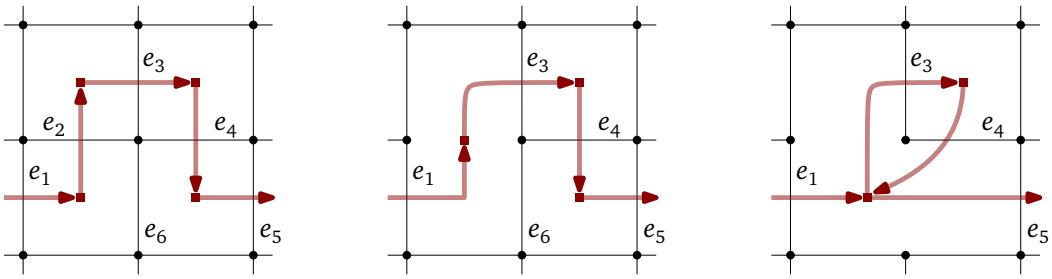

Figure 12: Transforming a cycle in G^* into a non-crossing closed walk in P^*. The portion $e_1^*, \ldots, e_5^*$ of a cycle in G^* (left) becomes $e_1^*, e_3^*, e_4^*, e_5^*$ with the deletion of e_2 (center) and keeps being $e_1^*, e_3^*, e_4^*, e_5^*$ with the deletion of e_6 (right).

Lemma 13. *The abstract Voronoi diagram of any* 3 *sites in the outer face of G has at most one vertex, besides a_∞.*

Proof. Assume that S is the set of 3 sites. Since each site $s \in S$ is in the outer face, the abstract Voronoi diagram $\mathrm{AVD}(s, S)$ contains the dual face s^*, which is incident to a_∞. It follows that all faces have a common vertex in a_∞. Since a plane graph with 3 faces can have at most 2 vertices, the result follows. □

5.3 Dealing with holes

Let G be a plane graph and let P be a connected subgraph of G, with the embedding inherited from G. Consider the graphic Voronoi diagram in P *using the distances in* G. Thus, for a set of weighted sites S in P and a site $s \in S$ we are interested in the graphic Voronoi diagrams

$$\mathrm{GVR}_{P,G}(s, S) \;=\; \mathrm{GVR}_G(s, S) \cap V(P).$$

We would like to represent these Voronoi diagrams using the dual graph of P^*. In particular, we have to define bisectors using the graph P^*.

Lemma 14. *Given two sites s and t, there is a non-crossing closed walk γ in P^* such that* $\mathrm{GVR}_{P,G}(s, \{s, t\})$ *and* $\mathrm{GVR}_{P,G}(t, \{s, t\})$ *are precisely $V_{\mathrm{int}}(\gamma, P)$ and $V_{\mathrm{ext}}(\gamma, P)$. Moreover, γ is obtained from* $\mathrm{bis}_G(s, t)$ *by deleting the edges of $(E(G) \setminus E(P))^*$ from the sequence of edges defining* $\mathrm{bis}_G(s, t)$.

Proof. Let $e_1^*, \ldots, e_k^*$ be the sequence of edges of G^* that define $\mathrm{bis}_G(s, t)$. If in this sequence we delete all appearances of e^* for $e \in E(G) \setminus E(P)$, then we obtain a subsequence $(e_1')^*, \ldots, (e_\ell')^*$ that defines a closed walk γ in P^*. See Figure 12 for a small example and Figure 13 for a larger example. The resulting closed walk is non-crossing, as can be seen by induction on the number of deleted edges. Indeed, if a plane graph H' is obtained from a plane graph H by deleting an edge e, then $(H')^*$ is obtained from H^* by contracting e^*. Any non-crossing walk in H^* remains non-crossing when contracting the edge $e^* \in E(H^*)$.

Note that the description using dual edges is simpler than using dual vertices because the relevant faces may also change with deletions of edges that do no appear in the cycle. We need that P is connected because otherwise P has faces that are not simply-connected and some closed walks of G^* may become empty (no edge of P is crossed) and in simply connected faces we have options to decide how to route the curve within a face; it is not enough to know the edges where we enter an exit the face. □

We use $\mathrm{bis}_{P,G}(s, t)$ for the non-crossing closed γ in P^* defined by Lemma 14. To use abstract Voronoi diagrams we have the following technical problem: in general, the curve $\mathrm{bis}_{P,G}(s, t)$ is not simple. We can work around this symbolically, as follows. Combinatorially, we keep encoding

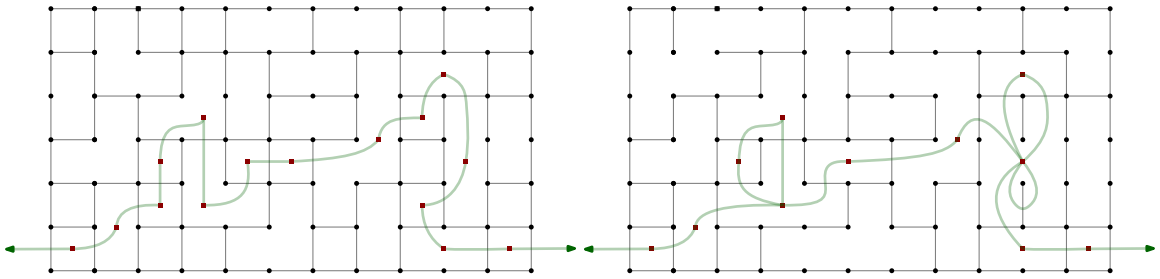

Figure 13: Transforming a cycle in G^* into a non-crossing closed walk in P^*.

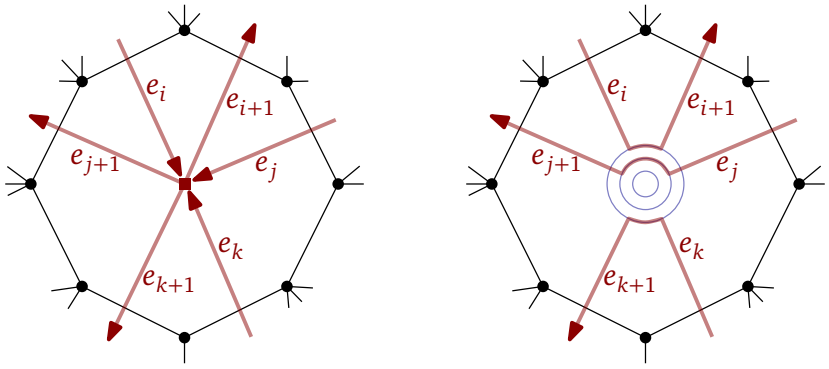

Figure 14: Transforming a closed walk in P^* into a simple curve.

the bisector as a closed walk in the dual graph P^*. However, the geometric curve associated to a description goes out of the dual graph to become simple. For each two consecutive edges aa' and $a'a''$ of each such closed walk, we always make a small shortcut in a small neighborhood of a' that avoids a'. For example, we can reroute the arcs along small concentric circles, where we use a larger radius when the distance along the face is smaller. See Figure 14 for an example. There are different ways to do this rerouting. In any case, the algorithm of Theorem 7 to build the abstract Voronoi diagram never uses coordinates. In such a way we obtain true geometric simple curves associated to each such bisector.

The transformation is not made for the outer face. Indeed, to use the technology of abstract Voronoi diagrams, we need that all the bisectors pass through a common point p_∞, which is a_∞. Thus, we do not want to make any rerouting at the outer face. This is not a problem if each bisector $\mathrm{bis}_{P,G}(s,t)$ passes exactly once through the vertex a_∞. If G and P have the same outer face, then $\mathrm{bis}_{P,G}(s,t)$ only passes once through a_∞. Thus, we will restrict attention to the case when G and P have the same outer face.

The rest of the presentation used for the case $G = P$ goes essentially unchanged. However, note that Lemma 8 does not hold in this case. An easier way to visualize things is to consider the creation of abstract Voronoi diagrams in the graph G^* and then consider the deletion of $E(G) \setminus E(P)$ in G (and in G^*). To summarize, we obtain the following.

Lemma 15. *Let G be a plane graph, let P be a connected subgraph of G such that G and P have the same outer face. Let S be a generic, independent set of sites located in the outer face of P. Let a_∞ be the vertex of G^* dual to the outer face of G. Then the system of abstract bisectors $\{(\mathrm{bis}_{P,G}(s,t), D_{P,G}(s,t)) \mid s,t \in S, s \neq t\}$ is admissible, where a_∞ plays the role of p_∞ in axiom (A3).*

The abstract Voronoi diagram of any 3 sites in the outer face of P has at most one vertex, besides a_∞.

For each $s \in S$, we have $\mathrm{GVR}_{P,G}(s,S) = V(P) \cap \mathrm{AVR}(s,S)$

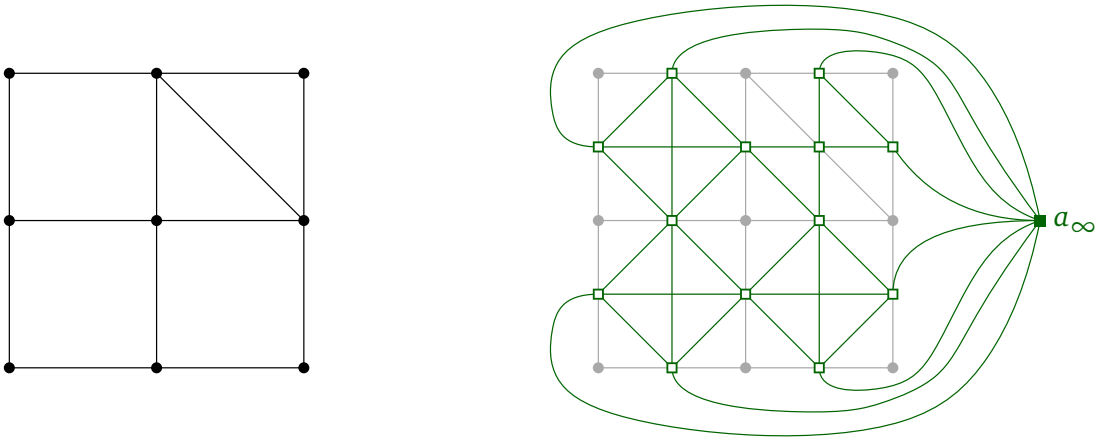

Figure 15: Left: a plane graph G. Right: a drawing of L_∞ in green (and G in gray).

Remark. Instead of using rerouting in the dual graph, another alternative is to use a variant of the line graph of the dual graph. The variant is designed to ensure that all the bisectors pass through a_∞, so that we can use abstract Voronoi diagrams. Let us spell out the construction of the graph, which we denote by L_∞. The vertex set of L_∞ is $E^*(G) \cup \{a_\infty\}$. Thus, each dual edge and the vertex a_∞ corresponding to the outer face of G are the vertices. For each face f of G that is not the outer face, we put an edge in L_∞ between each pair of edges that appear in f. For each edge e on the outer face of G, we put an edge in L_∞ between e^* and a_∞. This finishes the description of L_∞. The graph L_∞ has a natural drawing inherited from the embedding of G, that is not necessarily an embedding. (L_∞ has large cliques when G has large faces.) However, we can use a drawing of L_∞ to represent the curves. See Figure 15 for an example of (a drawing of) L_∞.

6 Algorithmic aspects of Voronoi diagrams in planar graphs

For the rest of this section, we assume that G is a connected plane graph, P is a connected subgraph of G, and the outer face of P and G coincide. We use r for the number of vertices in P. Let X be a set of b vertices in the outer face of P. We are interested in placing the sites at the vertices of X. In this section we assume that the distances $d_G(\cdot,\cdot)$ from each vertex of X to each vertex of P are known and available. We remark that the arcs of G may have negative weights, but G should not have negative cycles..

We next provide tools to manipulate portions of the bisectors and construct Voronoi diagrams in planar graphs.

Lemma 16. *For any two generic, independent sites $\{s,t\}$ placed at X we can compute $\mathrm{bis}_{P,G}(s,t)$ in $O(r)$ time.*

Proof. For each vertex $x \in V(P)$, we compare $w_s + d_G(s,x)$ and $w_t + d_G(t,x)$ to decide whether x belongs to $\mathrm{GVR}_{P,G}(s,\{s,t\})$ or $\mathrm{GVR}_{P,G}(t,\{s,t\})$. Note that $w_s + d_G(s,x) \neq w_t + d_G(t,x)$ because we assume generic sites. The sets $\mathrm{GVR}_{P,G}(s,\{s,t\})$ and $\mathrm{GVR}_{P,G}(t,\{s,t\})$ are nonempty because we assume independent sites. Now we can mark the edges of P with one endpoint in each of those sets and construct the closed walk $\mathrm{bis}_{P,G}(s,t)$ using the dual graph. □

Lemma 17. *Consider any two vertices $\{v_s, v_t\} \subseteq X$ as placements of sites. Consider the family of bisectors $\mathrm{bis}_{P,G}((v_s,w_s),(v_t,w_t))$ as a function of the weights w_s and w_t. There are at most $O(r)$ different bisectors. We can compute and store all the bisectors in $O(r^2)$ time such that, given two*

values w_s and w_t, the corresponding representation of $\mathrm{bis}_{P,G}((v_s,w_s),(v_t,w_t))$ *is accessed in* $O(\log r)$ *time.*

Proof. From the definition it is clear that

$$\mathrm{bis}_{P,G}((v_s,w_s),(v_t,w_t)) = \mathrm{bis}_{P,G}((v_s,0),(v_t,w_t-w_s)).$$

Thus, it is enough to consider the bisectors $\mathrm{bis}_{P,G}((v_s,0),(v_t,w))$ parameterized by $w\in\mathbb{R}$. Each bisector $\mathrm{bis}_G((v_s,0),(v_t,w))$ is a cycle in the dual graph G^* and the cycles are nested: as w increases, the graphic dominance region $\mathrm{GD}_G(s,t)$ monotonically grows and $D_G(s,t)$ also monotonically increases. The same happens with $\mathrm{bis}_{P,G}((v_s,0),(v_t,w))$: as w increases, the bisectors $\mathrm{bis}_{P,G}((v_s,0),(v_t,w))$ are nested and the region on one side monotonically grows. Since any two different non-crossing closed walks $\mathrm{bis}_{P,G}((v_s,0),(v_t,w))$ are nested and must differ by at least one vertex of P that is enclosed, there are at most $O(r)$ different bisectors.

For each vertex $x\in V(P)$, define the value $\eta_x = d_G(s,x)-d_G(t,x)$. The vertex x is in $\mathrm{GD}_G(s,t)$ when $w<\eta_x$, in $\mathrm{GD}_G(t,s)$ when $w>\eta_x$, and we have a degenerate (non-generic) case when $w=\eta_x$. Thus, we can compute the values $\{\eta_x \mid x\in V(P)\}$, sort them and store them sorted in a table. For each w between two consecutive values of $\{\eta_x \mid x\in V(P)\}$ we compute the bisector using Lemma 16 and store it with its predecessor of $\{\eta_x \mid x\in V(G)\}$. Given a query with shifts w_s,w_t, we use binary search in $O(\log r)$ time for the value w_t-w_s and locate the relevant bisector. □

As mentioned before, an abstract Voronoi vertex is just a vertex of P^* and an abstract Voronoi edge is encoded in the dual graph P^* by a tuple (s,t,aa',bb'), meaning that the edge is the portion of $\mathrm{bis}_{P,G}(s,t)$ starting with the dual edge aa' and finishing with the the dual edge bb' in some prescribed order, like for example the clockwise order of $\mathrm{bis}_{P,G}(s,t)$.

Lemma 18. *Consider any three vertices $\{v_q,v_s,v_t\}\subseteq X$ as placements of sites. Consider the family of abstract Voronoi diagrams for the sites $r=(v_q,w_q)$, $s=(v_s,w_s)$, and $t=(v_t,w_t)$ as a function of the weights w_q, w_s and w_t. We can compute and store all those Voronoi diagrams in $O(r^2)$ time such that, given the values w_q, w_s and w_t, the corresponding representation of the abstract Voronoi diagram of those 3 sites is accessed in $O(\log r)$ time.*

Proof. We use Lemma 17 to compute and store all the possible bisectors of each pair of vertices. This takes $O(r^2)$ time because we have $O(1)$ pairs of placements.

Only the difference between weights of the sites is relevant. Thus, we can just assume that the weight w_q is always 0. The relevant abstract Voronoi diagrams can thus be parameterized by the plane $\mathbb{R}^2$. The first coordinate is the weight w_s and the second coordinate is the weight w_t.

For each vertex $x\in V(P)$, we compute

$$\eta_x^{qs} = d_G(q,x)-d_G(s,x),\quad \eta_x^{qt} = d_G(q,x)-d_G(t,x),\quad \eta_x^{st} = d_G(s,x)-d_G(t,x).$$

Note that, once we fix the weights w_s,w_t and $w_q=0$, the vertex $x\in V(P)$ belongs to $\mathrm{GVR}_{P,G}(s,\{q,s,t\})$ if and only if $w_s<\eta_x^{qs}$ and $w_t-w_s>\eta_x^{st}$. A similar statement holds for the other sites, q and t.

In the plane (w_s,w_t) we consider the set of lines L that contains precisely the following lines

$$\begin{aligned}
&\forall x\in V(G):\ \ell_x^{qs} = \{(w_s,w_t)\in\mathbb{R}^2 \mid w_s=\eta_x^{qs}\},\\
&\forall x\in V(G):\ \ell_x^{qt} = \{(w_s,w_t)\in\mathbb{R}^2 \mid w_t=\eta_x^{qt}\},\\
&\forall x\in V(G):\ \ell_x^{st} = \{(w_s,w_t)\in\mathbb{R}^2 \mid w_t-w_s=\eta_x^{st}\}.
\end{aligned}$$

Since L has $O(r)$ lines, it breaks the plane $\mathbb{R}^2$ into $O(r^2)$ cells, usually called the arrangement induced by L and denoted by $\mathscr{A}(L)$. Such an arrangement can be computed in $O(r^2)$ time [11, Section 8.3]. For each cell $c\in\mathscr{A}(L)$, the Voronoi diagram defined by the sites $\{(q,0),(s,w_s),(t,w_t)\}$ is the same for all $(w_s,w_t)\in c$.

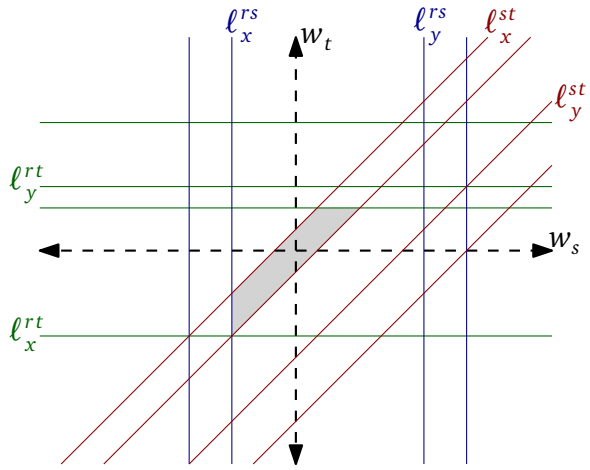

Figure 16: Example of lines L defined by 4 vertices. The cell of $\mathscr{A}(L)$ containing the origin is marked in gray. From the labels in th figure we can deduce that, in the Voronoi diagram of $q=(v_q,0)$, $s=(v_s,0)$ and $t=(v_t,0)$, the vertex x belongs to $\mathrm{GVR}_{P,G}(q,\{q,s,t\})$ and the vertex y belongs to $\mathrm{GVR}_{P,G}(s,\{q,s,t\})$.

We can further preprocess $\mathscr{A}(L)$ for standard point location [40]. Thus, after $O(r^2)$ preprocessing, given a query point (w_s,w_t), we can identify in $O(\log r)$ time the cell of $\mathscr{A}(L)$ that contains it.

In each cell c of $\mathscr{A}(L)$ we store a description of the Voronoi diagram defined for weights on that cell. We can compute the relevant Voronoi diagram for each cell in $O(1)$ amortized time using a traversal of $\mathscr{A}(L)$. A simple way is as follows. Consider any line ℓ of L. Let us say that $\ell=\ell_x^{qs}\in L$; the other cases are similar. Let ℓ_ε be an upward shift of ℓ by an infinitesimal $\varepsilon>0$. The value w_s remains constantly equal to $\eta_x^{qs}+\varepsilon$ as we walk along ℓ_ε, while the value w_t changes. Consider the bisector $\mathrm{bis}_{P,G}((v_q,0),(v_s,w_s))$ and let $e_1,\dots,e_k$ be the edges of P that it crosses, as we walk from a_∞ to a_∞. Thus, the bisector is actually the non-crossing closed walk $e_1^*,\dots,e_k^*$ in the dual graph P^*. For each such edge e_i, can compute a value $\zeta(e_i)$ such that e_i is part of the abstract Voronoi edge of $\{(v_q,0),(v_s,w_s),(v_t,w_t)\}$ that separates the cell of q and s if and only $w_t>\zeta(e_i)$. Indeed, if y_q is the endpoint of e_i closer to q and y_s the other endpoint, then e_i is (part of) an abstract Voronoi edge of $\{(v_q,0),(v_s,w_s),(v_t,w_t)\}$ if and only if

$$w_q+d_G(q,y_q)<w_t+d_G(t,y_q)\quad\text{and}\quad w_s+d_G(s,y_s)<w_t+d_G(t,y_s).$$

Using that $w_q=0$ and $w_s=\eta_x^{qs}+\varepsilon$, this is equivalent to the condition

$$w_t\;>\;\max\{d_G(q,y_q)-d_G(t,y_q),\;d_G(s,y_s)-d_G(t,y_s)+\eta_x^{qs}+\varepsilon\}\;=:\;\zeta(e_i).$$

Because of planarity, the values $\zeta(e_1),\dots,\zeta(e_k)$ are either monotonically increasing or decreasing. Indeed, the cell for t can only grow when w_t increases. Therefore, the values $\zeta(e_1),\dots,\zeta(e_k)$ are obtained already sorted. As we walk along ℓ_ε, we can identify the last edge e_i such that $\zeta(e_i)<w_t$ and identify the precise portion of $\mathrm{bis}_{P,G}((q,0),(s,w_s))$ that is in the Voronoi diagram of $\{(q,0),(s,w_s),(t,w_t)\}$.

Repeating this procedure for each line $\ell_x^{qs}\in L$, with two infinitesimal shifts per line, one on each side, we can figure out in $O(1)$ amortized time per cell the portion of $\mathrm{bis}_{P,G}(q,s)$ in the abstract Voronoi diagram for each cell of $\mathscr{A}(L)$ bounded by one of those lines. If a cell is not bounded by a line ℓ_x^{qs} for some x, we figure out this information from a neighbour cell. A similar approach for the lines $\ell_x^{qt}\in L$ and $\ell_x^{st}\in L$ determines the portions of $\mathrm{bis}_{P,G}(q,t)$ and $\mathrm{bis}_{P,G}(s,t)$, respectively. Thus, we obtain the abstract Voronoi diagrams for all cells $c\in\mathscr{A}(L)$ in $O(1)$ amortized time per cell. □

Recall that b is the cardinality of X.

Lemma 19. *There is a data structure with the following properties. The preprocessing time is $O(b^3r^2)$. For any generic, independent set S of 4 sites placed on X, the abstract Voronoi diagram $\mathrm{AVD}(S)$ can be computed in $O(\log r)$ time. The output is given combinatorially as a collection of abstract Voronoi vertices and edges encoded in the dual graph P^*.*

Proof. First, we make a table $T_X[\cdot]$ such that, for $u \in X$, $T_X[u]$ is the rank of u when walking along the boundary of the outer face of P and, for $u \notin X$, we have $T_X[u]$ undefined. Thus, given 3 vertices of X we can deduce their circular ordering along the boundary of the outer face of P in $O(1)$ time.

We use Lemma 17 to compute and store all the possible bisectors. Since there are b^2 different possible locations for the sites, for each pair of locations there are $O(r)$ different bisectors, and for each bisector we spend $O(r)$ space and preprocessing time, we have spent a total of $O(b^2r^2)$ time.

For each bisector β, we preprocess it to quickly figure out the circular order of its (dual) edges: given two edges aa' and bb' on β, is the clockwise order along β given by aa', bb', a_∞ or by bb', aa', a_∞? For each bisector β we can make a table $T_\beta[\cdot]$ indexed by the edges such that $T_\beta[aa']$ is the position of aa' along β, when we walk β clockwise starting from a_∞. We set $T_\beta[aa']$ to undefined when aa' does not appear in β. Thus, given 2 edges of β, we can decide their relative order along β in $O(1)$ time. The time and space for this, over all bisectors, is also $O(b^2r^2)$.

We make a table indexed by triples of vertices of X and, for each triple, we use Lemma 18 and store in the table a pointer to the resulting data structure. We have $O(|X|^3) = O(b^3)$ choices for the vertices hosting the sites, and thus we spend $O(b^3r^2)$ in the preprocessing step. Given any three sites placed at X, we can get the abstract Voronoi diagram of those three sites in $O(\log r)$ time. This finishes the preprocessing.

Assume that we are given a set S of 4 sites placed at X and we want to compute its abstract Voronoi diagram. We recover the abstract Voronoi diagrams for each subset $\binom{S}{3}$ in $O(\log r)$ time, using the stored data.

If there are two sites $s, t \in S$ such that their bisector $\mathrm{bis}_{P,G}(s,t)$ is in full in the Voronoi diagram of each subset S' with $|S'| = 3$ and $\{s,t\} \subset S' \subset S$, then in the abstract Voronoi diagram of S there is a region bounded only by $\mathrm{bis}_{P,G}(s,t)$. We can then compose that bisector and the abstract Voronoi diagram of the other three sites to obtain the final Voronoi diagram. See the left of Figure 17. (It may be that we have more than one such "isolated" abstract Voronoi region.)

In the opposite case, in the abstract Voronoi diagram there is no abstract Voronoi region that is bounded by a unique bisector. The abstract Voronoi diagram restricted to the interior faces of G is connected. The shape of such a Voronoi diagram can be only one of two, depending on which opposite sites share a common edge. See the center and right side of Figure 17. Let p, q, s, t be the sites in clockwise order along the boundary of G. We can infer this order in $O(1)$ time through the table $T_X[\cdot]$. Assume, by renaming the sites if needed, that $\mathrm{bis}_{P,G}(s,p)$ has s in its interior. From $\mathrm{AVD}(\{p,q,s\})$ we obtain the edge aa' of $\mathrm{bis}_{P,G}(s,p)$ incident to the vertex of $\mathrm{AVD}(\{p,q,s\})$, and from $\mathrm{AVD}(\{p,s,t\})$ we obtain the edge bb' of $\mathrm{bis}_{P,G}(s,p)$ incident to the vertex of $\mathrm{AVD}(\{p,s,t\})$. If $a = b$, then $\mathrm{AVD}(\{p,q,s,t\})$ has a common vertex of degree 4 that is incident to four abstract Voronoi edges. If $aa' = bb'$ or the the cyclic order of a_∞, bb', aa' along $\mathrm{bis}_{P,G}(p,s)$ is clockwise, then the tuple (p,s,aa',bb') defines an abstract edge in the abstract Voronoi diagram of S. Otherwise, there is tuple (q,t,cc',dd') for some edges cc' and dd' that can be obtained by exchanging the roles of p,s with q,t. From this information and the abstract Voronoi diagrams of each three sites, we can construct the abstract Voronoi diagram of $S = \{p,q,s,t\}$. □

Theorem 20. *Let G be a plane graph and let P be a connected subgraph of G with r vertices such that G and P have the same outer face. Let X be a set of b vertices on the outer face of P. Assume*

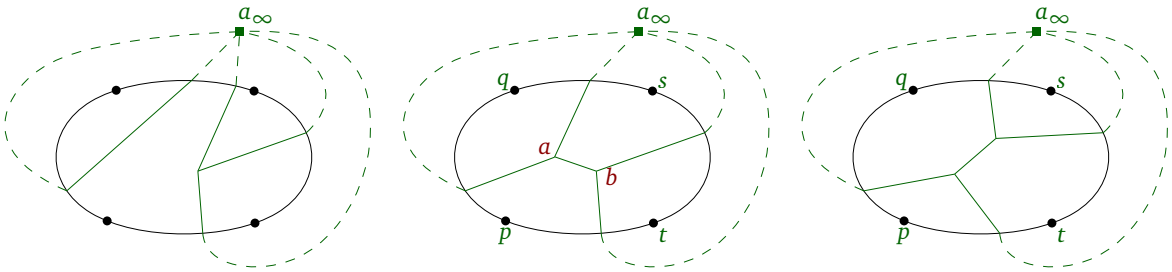

Figure 17: Left: an abstract Voronoi region is bounded by a single bisector. Center and right: possible configurations of the abstract Voronoi diagram of 4 sites, when it is connected.

that G has no negative cycles and the distances $d_G(\cdot,\cdot)$ from each vertex of X to each vertex of P are available. There is a data structure with the following properties. The preprocessing time is $O(b^3r^2)$. For any generic and independent set S of sites placed at X, the abstract Voronoi diagram $\mathrm{AVD}_{P,G}(S)$ can be computed in $\tilde{O}(b)$ expected time. The output is given combinatorially as a collection of abstract Voronoi vertices and edges encoded in the dual graph P^*.

Proof. We apply the preprocessing of Lemma 19. We spend $O(b^3r^2)$ time and, given any four sites placed on X, we can compute its abstract Voronoi diagram in $O(\log r)$ time.

Assume that we are given a set S of b sites placed at vertices of X. Because of Lemma 15 (see also Lemma 13), any three sites have a vertex in common, besides the one at p_∞ (or a_∞). According to Theorem 7, we can compute the abstract Voronoi diagram using $O(|S|\log r) = O(b\log r)$ expected time and expected number of elementary operations, where an elementary operation is the computation of an abstract Voronoi diagram of 4 sites. Since each elementary operation takes $O(\log r)$ time because of the data structure of Lemma 19, the result follows. □

7 Data structure for planar graphs

In this section we are going to use abstract Voronoi diagrams and the data structures of Section 3 to compute information about the distances from a fixed vertex in a planar graph when the length of the edges incident to the fixed vertex are specified at query time.

Let G be a plane graph with n vertices and let P be a connected subgraph of G with r vertices such that G and P have the same outer face. Let X be a set of b vertices on the outer face of P. Let U be a subset of $V(P)$. The graph G may have arcs with negative edges, but it does not have any negative cycle.

For each subset $Y \subset X$, let $G^+(Y)$ be the graph obtained from G by adding a new vertex x_0 and arcs $E_0(Y) = \{x_0{\to}y \mid y \in Y\}$. See Figure 18. We want to preprocess G and P for different types of queries, as follows. At preprocessing time, the lengths of the edges in $E_0(X)$ are undefined, unknown. At query time we are given a subset $Y \subseteq X$ and the lengths $\lambda(x_0{\to}y)$ for the arcs $x_0{\to}y$ of $E_0(Y)$. Using the notation introduced in Section 2, we are interested in the following information about the distances from the new vertex x_0:

$$\begin{aligned}
\mathrm{diam}(x_0, U, G^+(Y)) &= \max\{d_{G^+(Y)}(x_0,u) \mid u \in U\},\\
\mathrm{sum}(x_0, U, G^+(Y)) &= \sum_{u\in U} d_{G^+(Y)}(x_0,u),\\
\mathrm{count}_{\le}(x_0, U, G^+(Y), \delta) &= |\{u \in U \mid d_{G^+(Y)}(x_0,u) \le \delta\}|.
\end{aligned}$$

Note that we are only using the distances to the subset $U \subseteq V(P)$.

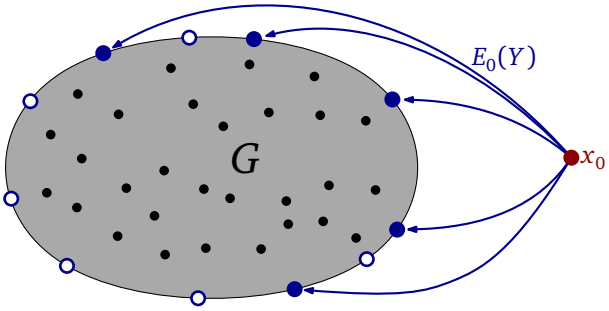

Figure 18: The graph $G^+(Y)$. The graph G is represented by the gray region, vertices of Y are marked with filled in dots, and vertices of $X \setminus Y$ are represented with void dots.

The set of vertices Y and the lengths $\lambda(x_0 \to y)$, where $y \in Y$, will be given so that they satisfy the following condition:

$$\forall y, y' \in Y,\ y \neq y' : \quad \lambda(x_0 \to y) < \lambda(x_0 \to y') + d_G(y', y). \tag{6}$$

This condition implies that, for all $y \in Y$, there is a unique shortest from x_0 to $y \in Y$ and this shortest path is just the arc $x_0 \to y$. This condition is important in our scenario to ensure that, when using the vertices of Y as sites with weights $\lambda(x_0 \to y)$, the sites are generic and independent.

Theorem 21. *Assume that G is a weighted plane graph with n vertices and no negative cycles. Let P be a subgraph of G with r vertices such that G and P have the same outer face. Let X be a set of b vertices on the outer face of P and let U be a subset of $V(P)$. After $\tilde{O}(n + b^3 r^2)$ preprocessing time, we can handle the following queries in $\tilde{O}(b)$ expected time: given a subset of vertices $Y \subset X$ and weights for the darts $\lambda(x_0 \to y)$, $y \in Y$, that satisfy the condition (6), return* $\mathrm{sum}(x_0, U, G^+(Y))$.

Proof. We compute and store the distances $d_G(x, v)$ for all $x \in X$ and $v \in V(P)$. This can be done in $\tilde{O}(n + br)$ time, as follows. First we compute a single-source shortest-path tree in $\tilde{O}(r)$ time[14, 25, 35]. With this we have a potential function in G and for the next distances we can assume non-negative weights. Then we use that all the vertices of X are incident to the outer face of G. Using [7, 23] we obtain in $\tilde{O}(|V(P)| + |X| \cdot |V(P)|) = \tilde{O}(n + br)$ time the distances $d_G(x, v)$, for all $x \in X$ and $v \in V(P)$.

We preprocess the pair of graphs G and P as described in Theorem 20. This takes $O(b^3 r^2)$ time because we already have all the required distances.

For each vertex $x \in X$ we proceed as follows. We compute all the bisectors of the type $\mathrm{bis}_{P,G}((x, \cdot), (x', \cdot))$ for all $x' \in X$. Let Π_x be the resulting family of curves. Then, we preprocess P with respect to Π_x as explained in Theorems 2. More precisely, we use Theorem 2 for the following two vertex-weight functions:

$$\omega_x(v) = \begin{cases} d_G(x, v), & \text{if } v \in U, \\ 0, & \text{otherwise.} \end{cases} \qquad \omega'_x(v) = \begin{cases} 1, & \text{if } v \in U, \\ 0, & \text{otherwise.} \end{cases}$$

We denote by $\sigma_x(\cdot)$ and $\sigma'_x(\cdot)$ the corresponding sums of weights. For example, $\sigma'_x(U') = \sum_{v \in U'} \omega'_x(v)$ for all $V' \subseteq V(P)$. This finishes the description of the preprocessing.

Let us analize the running time for the last step of the preprocessing. For each two vertices $x, x' \in X$, we spend $O(r^2)$ time to compute the bisectors $\mathrm{bis}_{P,G}((x, \cdot), (x', \cdot))$ because of Lemma 17. It follows that Π_x is a family of walks in P^* with $O(br^2)$ dual edges, counted with multiplicity. The preprocessing of Theorem 2 is $O(r + ||\Pi_x||) = O(r + br^2) = O(br^2)$ per vertex $x \in X$, where $||\Pi_x||$ denotes the number of edges in Π_x. Thus, over all $x \in X$, we spend $O(b^2 r^2)$ time.

Consider now a query specified by a subset $Y \subset X$ and the edge weights $\lambda(x_0 \to y)$, $y \in Y$, that satisfy the condition (6). For each $y \in Y$, define the site $s_y = (y, \lambda(x_0 \to y))$. Because of

condition (6), the set of sites $S_Y = \{s_y \mid y \in Y\}$ is independent and generic. Using the data structure of Theorem 20, we compute the weighted Voronoi diagram for the sites S_Y. Thus, we obtain the abstract Voronoi diagram in $\tilde{O}(|Y|) = \tilde{O}(b)$ expected time. For each $y \in Y$, let γ_y be the closed walk in the dual graph P^* that defines the boundary of $\mathrm{AVR}(s_{y(v)}, S)$.

For each vertex $v \in V(P)$ there is precisely one vertex $y(v) \in Y$ such that $\mathrm{AVR}(s_{y(v)}, S)$ contains v. Moreover, because of the definition of (graphic) Voronoi diagrams, we have

$$d_{G^+(Y)}(x_0, v) \;=\; \lambda(x_0 \to y(v)) + d_G(y(v), v).$$

Note that

$$\begin{aligned}
\mathrm{sum}(x_0, U, G^+(Y)) \;&=\; \sum_{u \in U} d_{G^+(Y)}(x_0, u) \\
&= \sum_{u \in U} (\lambda(x_0 \to y(v)) + d_G(y(v), v)) \\
&= \sum_{y \in Y} \sum_{u \in U \text{ s.t.} y(u) = y} (\lambda(x_0 \to y(v)) + d_G(y(v), v)) \\
&= \sum_{y \in Y} \sum_{u \in \mathrm{AVR}(s_y, S) \cap U} (\lambda(x_0 \to y(v)) + d_G(y, v)) \\
&= \sum_{y \in Y} \left(\lambda(x_0 \to y(v)) \cdot |\mathrm{AVR}(s_y, S) \cap U| + \sum_{u \in \mathrm{AVR}(s_y, S) \cap U} d_G(y, v) \right) \\
&= \sum_{y \in Y} \Big(\lambda(x_0 \to y(v)) \cdot \sigma'_y(V_{\mathrm{int}}(\gamma_y, P)) + \sigma_y(V_{\mathrm{int}}(\gamma_y, P)) \Big) \qquad (7)
\end{aligned}$$

For each site $y \in Y$, we walk along γ_y, the boundary of the abstract Voronoi region $\mathrm{AVR}(s_y, Y)$, and use the data structures of Theorem 2 for ω_y and ω'_y to collect the data

$$\forall y \in Y: \quad \sigma_y(V_{\mathrm{int}}(\gamma_y, P)), \text{ and } \sigma'_y(V_{\mathrm{int}}(\gamma_y, P)).$$

Here we are using that $y \in \mathrm{AVR}(s_y, S)$, and thus y is in the interior of γ_y. For each γ_y we spend $\tilde{O}(1)$ times the complexity of its description. Over all Y, this takes $\tilde{O}(|Y|) = \tilde{O}(b)$ time. From this information we can compute $\mathrm{sum}(x_0, U, G^+(Y))$ using (7), and the result follows. □

Theorem 22. *Consider the setting of Theorem 21. After $\tilde{O}(nb + b^3r^2 + b^4)$ preprocessing time, we can handle the following queries in $\tilde{O}(b)$ expected time: given a subset of vertices $Y \subset X$, weights for the darts $\lambda(x_0 \to y)$, $y \in Y$, that satisfy the condition (6), return* $\mathrm{diam}(x_0, U, G^+(Y))$.

Proof. We use the same approach as in the proof of Theorem 21. We keep using the notation of that proof. The main difference is that we do not use the data structure of Theorem 2, but the data structure of Theorem 5. We explain the details of this part.

For each vertex $x \in X$ we proceed as follows. Let T_x be a shortest-path tree in G from x. We do not compute T_x, but use it to argue correctness. Then, we use the data structure of Theorem 5 for G, P, the tree T_x, and the vertex-weights $\omega_x(\cdot)$. Let μ_x be the corresponding maximum function that returns the data structure. Thus, $\mu_x(U) = \max\{\omega_x(u) \mid u \in U\}$. This finishes the description of the preprocessing.

Let us analize the running time for this step of the preprocessing. Like before, each Π_x is computed in $O(br^2)$ time and has $O(br^2)$ dual edges, counted with multiplicity. The preprocessing of Theorem 5 is $O(n + ||\Pi_x|| + b^3) = O(n + br^2 + b^3)$ time for each $x \in X$. Therefore, the total preprocessing used in this step is $O(nb + b^2r^2 + b^4)$.

Next, we note that each γ_y is in $\widetilde{\Xi}(G, P, T_y)$ because of Lemma 8. Therefore, we can obtain $\mu_y(V_{\text{int}}(\gamma_y, P))$ in $\tilde{O}(1)$ times the complexity of the description of γ_y. Over all $y \in Y$, this takes $\tilde{O}(|Y|) = \tilde{O}(b)$ time.

With this data, the desired value is then obtained in $O(|Y|) = O(b)$ time using that

$$\begin{aligned}
\operatorname{diam}(x_0, U, G^+(Y)) &= \max\{d_{G^+(Y)}(x_0, u) \mid u \in U\} \\
&= \max_{y \in Y} \max_{u \in U \text{ s.t.} y(u)=y} (\lambda(x_0 \to y(v)) + d_G(y(v), v)) \\
&= \max_{y \in Y} \left(\lambda(x_0 \to y(v)) + \max_{u \in U \text{ s.t.} y(u)=y} d_G(y(v), v) \right) \\
&= \max_{y \in Y} \left(\lambda(x_0 \to y(v)) + \mu_y(V_{\text{int}}(\gamma_y, P)) \right). \qquad \square
\end{aligned}$$

Theorem 23. *Consider the setting of Theorem 21. After $\tilde{O}(n + b^3 r^3)$ preprocessing time, we can handle the following queries in $\tilde{O}(b)$ expected time: given a subset of vertices $Y \subset X$, weights for the darts $\lambda(x_0 \to y)$, $y \in Y$, that satisfy the condition (6), and a real value δ, return $\operatorname{count}_{\leq}(x_0, U, G^+(Y), \delta)$.*

Proof. We use the same approach as in the proof of Theorem 21 and keep using its notation. The main difference is that we do not use the data structure of Theorem 2, but the data structure of Corollary 3 for the vertex-weights $\omega_x(\cdot)$. Let $\kappa^x_{\leq}$ be the corresponding function. This means that, for each $x \in X$, we spend an extra factor r in the preprocessing. Thus, for each x we spend $O(b^2 r^3)$ time, instead of $O(b^2 r^2)$. Over all $x \in X$, this means that the preprocessing has an extra factor of $O(b^3 r^3)$.

The rest of the approach is the same. We just have to use that

$$\begin{aligned}
\operatorname{count}_{\leq}(x_0, U, G^+(Y), \delta) &= |\{u \in U \mid d_{G^+(Y)}(x_0, u) \leq \delta\}| \\
&= \sum_{y \in Y} |\{u \in V_{\text{int}}(\gamma_y, P) \cap U \mid \lambda(x_0 \to y(v)) + d_G(y(v), v) \leq \delta\}| \\
&= \sum_{y \in Y} \kappa^y_{\leq}(V_{\text{int}}(\gamma_y, P), \delta - \lambda(x_0 \to y(v)),
\end{aligned}$$

and all values $\kappa^y_{\leq}(V_{\text{int}}(\gamma_y, P), \delta - \lambda(x_0 \to y(v))$ are recovered from the data structure of Corollary 3 in $\tilde{O}(|Y|) = \tilde{O}(b)$ time. $\square$

8 Diameter and Sum of Distances in Planar Graphs

The data structures of Theorems 21, 22 and 23 are going to be used for each piece of an r-division. Then, for each vertex of G we are going to query it. We first explain the precise concept of piece and division that we use, and then explain its use.

Divisions. The concept of r-division for planar graphs was introduced by Frederickson [16], and then refined and used by several authors; see for example [5, 19, 26, 36] for a sample. For us it is most convenient to use the construction of Klein, Mozes and Sommer [24]. We first state the definitions carefully, almost verbatim from [24].

Let G be a plane graph. A ***piece***[2] P of G is an edge-induced subgraph of G. In each piece we assume the embedding inherited from G. A ***boundary vertex*** of a piece P is a vertex of P that is incident to some edge in $E(G) \setminus E(P)$. A ***hole*** of a piece P is a face of P that is not a face of G. Note that all boundary vertices of a piece P are incident to a hole of P. An ***r-division with a few holes*** of G is a collection $\{P_1, \ldots, P_k\}$ of pieces of G such that

[2]They use the term "region", which in our opinion is more suitable. However, we are using such a term for so many things in this paper that in our context we prefer to use some other term.

- there are $O(n/r)$ pieces, that is, $k = O(n/r)$;
- each edge of G is in at least one piece;
- each piece has $O(r)$ vertices;
- each piece has $O(\sqrt{r})$ boundary vertices;
- each piece has $O(1)$ holes.

Theorem 24 (Klein, Mozes and Sommer [24])**.** *There is a linear-time algorithm that, for any biconnected triangulated planar embedded graph G, outputs an r-division of G with few holes.*

In fact, we will only use that all pieces together have $O(n/r)$ holes. Thus, other decompositions proposed by other authors could also be used. Note that we can assume that each piece is connected because we could replace each piece by its connected components, and we would get a new r-division with a few holes.

Work per piece. We now describe how to compute the relevant information within a fixed piece and the information between a fixed piece and all vertices outside the piece. The next result is enough for our usage; stronger can be obtained using, for example [33, 34]

Lemma 25. *Let P be a piece of G with r vertices and $O(\sqrt{r})$ boundary vertices. Let U be a subset of vertices in P. In $\tilde{O}(nr^{1/2}+r^2)$ time we can compute for all vertices $v \in V(P)$ the values* $\mathrm{diam}(v, U, G)$*,* $\mathrm{sum}(v, U, G)$*, and* $\mathrm{count}_{\leq}(v, U, G, \delta)$ *(for a given $\delta \in \mathbb{R}$).*

Proof. Let ∂ be the set of boundary vertices of P in G. We compute shortest-path trees from each vertex $x \in \partial$ in G in near-linear time [14, 25, 35]. This takes $|\partial| \cdot \tilde{O}(n) = \tilde{O}(nr^{1/2})$ time.

We build a graph $\widetilde{P}$ by adding to P arcs between each pair of vertices of ∂. The length of each new arc $x \to y$ is set to $d_G(x, y)$. Standard arguments show that a distance between any two vertices of P is the same in G and in $\widetilde{P}$. The graph $\widetilde{P}$ has $O(|E(P)| + |\partial|^2) = O(r)$ edges and $O(r)$ vertices. We can compute all pairwise distances in $\tilde{O}(|V(\widetilde{P})| \cdot |E(\widetilde{P})|) = O(r \cdot r) = O(r^2)$ time using standard approaches. (Since $\widetilde{P}$ may have negative weights, we may have to use a potential function.)

From all the distances in $\widetilde{P}$, that are also distances in G, we can compute the desired values directly. □

Lemma 26. *Let P be a piece of G with r vertices, $O(\sqrt{r})$ boundary vertices, and h holes. Let U be a subset of vertices in P.*

- *In $\tilde{O}(nh + r^{7/2} + nr^{1/2})$ expected time we can compute the values* $\mathrm{sum}(v, U, G)$ *for all vertices $v \in V(G) \setminus V(P)$.*
- *In $\tilde{O}(nh + r^{7/2} + nr^{1/2})$ expected time we can compute the values* $\mathrm{diam}(v, U, G)$ *for all vertices $v \in V(G) \setminus V(P)$.*
- *In $\tilde{O}(nh + r^{9/2} + nr^{1/2})$ expected time we can compute the values* $\mathrm{count}_{\leq}(v, U, G, \delta)$ *for a given δ and all vertices $v \in V(G) \setminus V(P)$.*

Proof. Let $C_1, \ldots, C_h$ be the facial walks of the holes of P. For $i \in [h]$, let A_i be the vertices of G contained in the hole defined by C_i. Since each vertex of $V(G) \setminus V(P)$ is contained in exactly one hole of P, the sets $A_1, \ldots, A_h$ form a partition of $V(G) \setminus V(P)$. For each $i \in [h]$, we define the graph G_i as the subgraph of G obtained by filling the holes of P, but the one that defines A_i. For each $i \in [h]$, let X_i be the set of boundary vertices that appear in C_i. See Figure 19 for an example. The sets $A_1, \ldots, A_h, X_1, \ldots, X_h$, and the graphs $G_1, \ldots, G_h$ can be constructed in $O(nh)$ time.

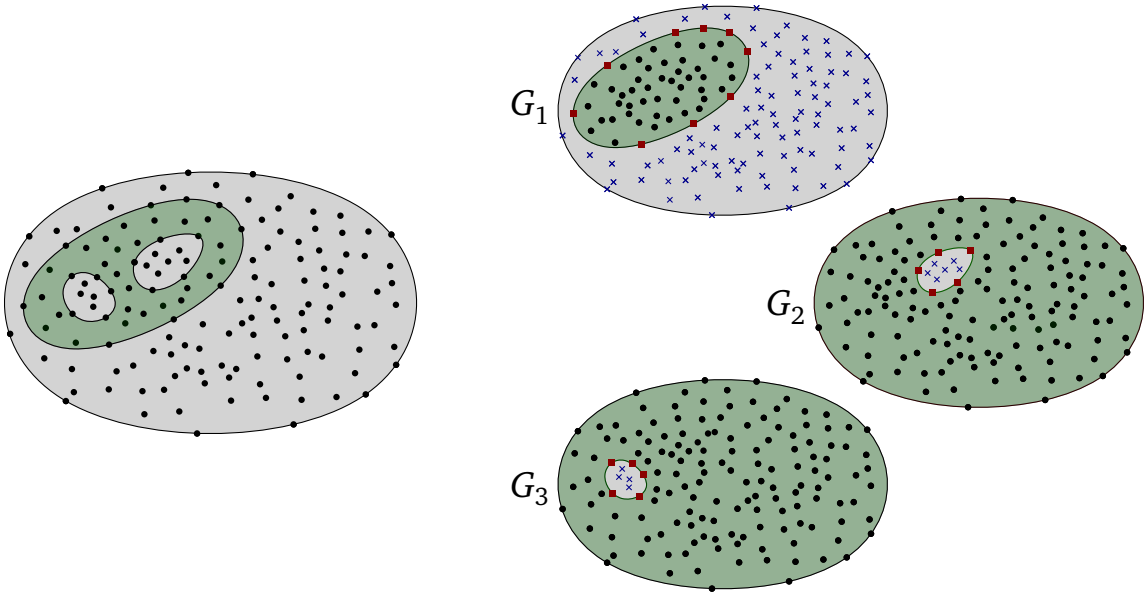

Figure 19: The graphs G_i and the vertices X_i in the proof of Lemma 26. On the left we have a graph G (in gray) and a piece P (in light green) of G with three holes. On the right we have, for $i \in [3]$, the graph G_i (in light green), the set X_i (red squares) and the set A_i (blue crosses).

We compute the distances in G and in G_i from and to all boundary vertices X_i. This can be done computing $4 \cdot |X_i| = O(\sqrt{r})$ different shortest-path trees, each of them in G, G_i, or the reversed graphs G^R, G_i^R. Since each single-source shortest path can be computed in $\tilde{O}(n)$ time [14, 25, 35], we spend in total $\tilde{O}(nr^{1/2})$ time.

Consider any fixed index $i \in [h]$. For each $v \in A_i$, let Y_i^v be the vertices y of X_i such that in the shortest path in G from v to y the last arc is not contained in G_i. For each $x \in X_i \setminus Y_i^v$, there exists some other boundary vertex $y \in Y_i^v$ such that $d_G(v,x) = d_G(v,y) + d_{G_i}(y,x)$. Therefore, for each $u \in V(P)$, we have

$$d_G(v,u) \;=\; \min\{d_G(v,x) + d_{G_i}(x,y) \mid x \in X_i\} \;=\; \min\{d_G(v,y) + d_{G_i}(y,u) \mid y \in Y_i^v\}.$$

Because the selection we made for Y_i^v and the uniqueness of shortest paths in G, we have that

$$\forall y, y' \in Y_i^v, y \neq y' : \quad d_G(v,y) < d_G(v,y') + d_{G_i}(y',y). \tag{8}$$

Using the shortest-path trees to $x \in X_i$, we can identify the relevant pairs $\{(v,x) \mid v \in A_i, x \in X_i^v\}$ in $O(n|X_i|)$ time. Since $\sum_i |X_i| = O(\sqrt{r})$, over all indices $i \in [h]$ we spend $O(nr^{1/2})$ time.

For each $i \in [h]$, fix an embedding of G_i such that C_i defines the outer face and thus X_i lies in the outer face of G_i and P.

Now there are slight differences depending on the data we want to compute. The difference lies in which data structure we use. Let us first consider the problem of computing $\mathrm{sum}(v, U, G)$. We apply Theorem 21 for the graph G_i and the piece P with respect to the set X_i. Since G_i has $O(n)$ vertices, P has $O(r)$ vertices, and X_i has $O(\sqrt{r})$ vertices, the preprocessing takes $\tilde{O}(n + |X_i|^3 r^2))$ time. Now, for each vertex $v \in A_i$, we consider the graph $G_i^+(Y_i^v)$ with edge weights $\lambda(x_0 \to y) = d_G(v,y)$ for all $y \in Y_i^v$. Note that, with these weights, the property in (8) corresponds to condition (6). Moreover, for each $u \in V(P)$ we have $d_G(v,u) = d_{G_i^+(X_i^v)}(x_0,u)$. Therefore, we can use the data structure of Theorem 21 to get in $\tilde{O}(|Y_i^v|) = \tilde{O}(|X_i|)$ time

$$\begin{aligned}\mathrm{sum}(x_0, U, G_i^+(Y_i^v)) \;&=\; \sum_{u \in U} d_{G_i^+(X_i^v)}(x_0,u) \;=\; \sum_{u \in U} d_G(v,u) \\ &=\; \mathrm{sum}(v, U, G).\end{aligned}$$

Iterating over all $i \in [h]$ and noting that $A_1, \ldots, A_h$ is a partition of $V(G) \setminus V(P)$, we obtain the desired values: $\text{sum}(v, U, G)$ for all $v \in V(G) \setminus V(P)$. The running time for the preprocessing of this last step is

$$\sum_i \tilde{O}(n + |X_i|^3 r^2) \;=\; \tilde{O}(nh + |X|^3 r^2) \;=\; \tilde{O}(nh + r^{7/2}).$$

and the running time for the queries is

$$\sum_i \tilde{O}(|A_i| \cdot |X_i|) \;=\; \tilde{O}\left((\sum_i |A_i|) \cdot (\sum_i |X_i|) \right) \;=\; \tilde{O}(nr^{1/2}).$$

The result in the first item follows.

For computing $\text{diam}(v, U, G)$, we use the same approach, but employ Theorem 22 instead of Theorem 21. The preprocessing time for $i \in [h]$ has an extra factor $\tilde{O}(|X_i|^4)$. Therefore, the preprocessing time in the last step becomes

$$\sum_i \tilde{O}(n + |X_i|^3 r^2 + |X_i|^4) \;=\; \tilde{O}(nh + |X|^3 r^2 + |X|^4) \;=\; \tilde{O}(nh + r^{7/2}).$$

The rest is essentially the same, and we obtain the claim in the second item.

For computing $\text{count}_{\leq}(v, U, G, \delta)$, we use the same approach, but employ Theorem 23 instead of Theorem 21. The preprocessing time for $i \in [h]$ has an extra factor $\tilde{O}(|X_i|^3 r^3)$. Therefore, the preprocessing time in the last step becomes

$$\sum_i \tilde{O}(n + |X_i|^3 r^3) \;=\; \tilde{O}(nh + |X|^3 r^3) \;=\; \tilde{O}(nh + r^{9/2}).$$

The rest is essentially the same, and we obtain the claim in the third item. □

Working over all pieces. We can now obtain our main result.

Theorem 27. *Let G be a planar graph with n vertices, real abstract length on its arcs, and no negative cycle. In $O(n^{11/6}\,\text{polylog}(n))$ expected time we can compute $\text{sum}(x, V(G), G)$ and $\text{diam}(x, V(G), G)$ for all vertices x of G. For a given $\delta \in \mathbb{R}$, in $O(n^{15/8}\,\text{polylog}(n))$ expected time we can compute $\text{count}_{\leq}(x, V(G), G, \delta)$ for all vertices x of G.*

Proof. Adding edges of sufficiently large lengths, we may assume that G is triangulated. We also embed G. These operations can be done in linear time. With a slight abuse of notation, we keep using G for the resulting embedded, triangulated graph.

We compute an r-division $\mathscr{P} = \{P_1, \ldots, P_k\}$ of G with few holes, for a parameter r to be specified below. According to Theorem 24, this takes $O(n)$ time.

To avoid double counting we assign each vertex to a unique piece, as follows. For each vertex x of G we select a unique index $i(x)$ such that the piece $P_{i(x)}$ contains x. For each piece $P_j \in \mathscr{P}$, we define the set $U_j = \{x \in V(P_j) \mid i(y) = j\}$. The sets $U_1, \ldots, U_k$ are a partition of $V(G)$ and can be easily computed in linear time.

Next, we iterate over the pieces and, for each piece $P_i \in \mathscr{P}$, we use Lemma 26 to compute the values

$$\text{sum}(v, U_i, G),\ \text{diam}(v, U_i, G) \quad \forall v \in V(G) \setminus V(P_i).$$

We also use Lemma 25 to compute

$$\text{sum}(v, U_i, G),\ \text{diam}(v, U_i, G) \quad \forall v \in V(P_i).$$

Since the piece P_i has $O(1)$ holes, we spend $\tilde{O}(n + r^{7/2} + nr^{1/2}) = \tilde{O}(r^{7/2} + nr^{1/2})$ time per piece. Iterating over the $O(n/r)$ pieces, we get

$$\mathrm{sum}_G(v, U_i),\ \mathrm{diam}_G(v, U_i), \quad \forall v \in V(G),\ i \in [k]$$

in time

$$O(n/r) \cdot \tilde{O}(r^{7/2} + nr^{1/2}) = \tilde{O}(nr^{5/2} + n^2 r^{-1/2}).$$

Because $U_1, \ldots, U_k$ is a partition of $V(G)$ we can easily compute the desired values because

$$\begin{aligned} \mathrm{sum}(v, V(G), G) &= \sum_{i \in [k]} \mathrm{sum}(v, U_i, G), \\ \mathrm{diam}(v, V(G), G) &= \max\{\mathrm{diam}(v, U_i, G) \mid i \in [k]\}. \end{aligned}$$

(For the diameter of course we do not need that the sets $U_1, \ldots, U_k$ are disjoint.) Taking $r = n^{1/3}$ the running time becomes $\tilde{O}(n^{11/6})$ in expectation.

For $\mathrm{count}_{\leq}(\cdot)$ we use the third item of Lemma 26 to compute for each piece P_i

$$\mathrm{count}_{\leq}(v, U_i, G, \delta) \quad \forall v \in V(G) \setminus V(P_i).$$

Then, for each piece we spend $\tilde{O}(n + r^{9/2} + nr^{1/2}) = \tilde{O}(r^{9/2} + nr^{1/2})$. Over all pieces, the running time thus becomes

$$O(n/r) \cdot \tilde{O}(r^{9/2} + nr^{1/2}) = \tilde{O}(nr^{7/2} + n^2 r^{-1/2}).$$

Choosing $r = n^{1/4}$ we obtain a running time of $\tilde{O}(n^{15/8})$. Again using that $U_1, \ldots, U_k$ is a partition of $V(G)$, we can compute

$$\mathrm{count}_{\leq}(v, V(G), G, \delta) = \sum_{i \in [k]} \mathrm{count}_{\leq}(v, U_i, G, \delta).$$

□

Corollary 28. *Let G be a planar graph with n vertices, real abstract length on its arcs, and no negative cycle. In $O(n^{11/6}\,\mathrm{polylog}(n))$ expected time we can compute the diameter of G and the sum of the pairwise distances in G. For a given $\delta \in \mathbb{R}$, in $O(n^{15/8}\,\mathrm{polylog}(n))$ expected time we can compute the number of pairs of vertices in G at distance at most δ.*

9 Discussion

We have decided to explain the construction through the use of abstract Voronoi diagrams, instead of providing an algorithm tailored to our case. It is not clear to author which option would be better. In any case, for people familiar with randomized incremental constructions, it should be clear that the details can be worked out, once the compact representation of the bisectors using the dual graph is available. Using a direct algorithm perhaps we could get rid of the assumption that the sites have to be in the outer face and perhaps we could actually build a deterministic algorithm.

There are deterministic algorithms to compute abstract Voronoi diagrams [27, 29]. However, they require additional elementary operations and properties. Also, when the abstract Voronoi diagram has a forest-like shape, it can be computed in linear time [3]. At this point, it is unclear to the author whether these results are applicable in our case.

We think that the algorithm can be extended to graphs on surfaces of small genus, but for this one should take care to extend the construction of abstract Voronoi diagrams to graphs on surfaces or to planar graphs when the sites are in $O(g)$ faces, where g is the genus of the surface.

Let us discuss the approach the reduction to the planar case. The first step is to find an r-division where each part is planar. For this we can use the separator theorem of Eppstein [13]. It computes a set of curves on the surface that pass through $O(\sqrt{gn})$ vertices of G, they do not pass through the interior of any edge, and cutting along them gives a collection of planar patches, possibly with multiple holes. Taking a maximal subset of the curves that are homologically independent, we will get $O(g)$ curves that pass through $O(\sqrt{gn})$ vertices and cut the surface into planar patches. Now we can compute an r-division in each of the patches. We can also compute the distances from all the boundary vertices because shortest paths in the presence of negative weights can be computed in subquadratic time [9, 31]. Now we run into problems. Consider a vertex x and a piece P. The shortest paths from x to different vertices of P can pass through different boundary cycles. In the planar case, we always have a boundary cycle that intersects all the paths from x to all vertices of P. There is no such cycle in the case of surfaces, for example, when one of the planar patches is a cylinder. Computing Voronoi diagrams for sites placed in $O(g)$ cycles would handle this problem. While we think that this should be doable, it requires non-trivial work. In particular, some of the holes may have non-trivial topology.

At this point, we think that the main problem is whether the exponent 11/6 can be substantially reduced. On the opposite direction, the author would be surprised if the problems considered in this paper can be solved in near-linear time. Thus, the author conjectures that there should be some conditional lower bounds of the type $\Omega(n^c)$ for some constant $c > 1$.

Acknowledgments

This work was initiated at the Dagstuhl seminar Algorithms for Optimization Problems in Planar Graphs, 2016. I am very grateful to Kyle Fox, Shay Mozes, Oren Weimann, and Christian Wulff-Nilsen for several discussions on the problems treated here.

References

[1] A. Abboud, V. Vassilevska Williams, and J. Wang. Approximation and fixed parameter subquadratic algorithms for radius and diameter in sparse graphs. *Proc. 27th ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*, pp. 377–391, 2016, http://dl.acm.org/citation.cfm?id=2884435.2884463.

[2] M. A. Bender, M. Farach-Colton, G. Pemmasani, S. Skiena, and P. Sumazin. Lowest common ancestors in trees and directed acyclic graphs. *J. Algorithms* 57(2):75–94, 2005, http://dx.doi.org/10.1016/j.jalgor.2005.08.001.

[3] C. Bohler, R. Klein, and C. Liu. Forest-like abstract voronoi diagrams in linear time. *Proc. 26th Canadian Conference on Computational Geometry, CCCG 2014*, 2014, http://www.cccg.ca/proceedings/2014.

[4] J. Bondy and U. Murty. *Graph theory*. Graduate texts in mathematics 244. Springer, 2008.

[5] S. Cabello. Many distances in planar graphs. *Algorithmica* 62(1-2):361–381, 2012, http://dx.doi.org/10.1007/s00453-010-9459-0.

[6] S. Cabello. Subquadratic algorithms for the diameter and the sum of pairwise distances in planar graphs. *Proc. 28th ACM-SIAM Symposium on Discrete Algorithms, SODA 2017*, pp. 2143–2152, 2017, http://dx.doi.org/10.1137/1.9781611974782.139.

[7] S. Cabello, E. W. Chambers, and J. Erickson. Multiple-source shortest paths in embedded graphs. *SIAM J. Comput.* 42(4):1542–1571, 2013, http://dx.doi.org/10.1137/120864271.

[8] S. Cabello and C. Knauer. Algorithms for graphs of bounded treewidth via orthogonal range searching. *Comput. Geom.* 42(9):815–824, 2009, http://dx.doi.org/10.1016/j.comgeo.2009.02.001.

[9] E. W. Chambers, J. Erickson, and A. Nayyeri. Homology flows, cohomology cuts. *SIAM J. Comput.* 41(6):1605–1634, 2012, http://dx.doi.org/10.1137/090766863.

[10] É. Colin de Verdière. Shortest cut graph of a surface with prescribed vertex set. *Algorithms - ESA 2010, 18th Annual European Symposium, Part II*, pp. 100–111. Springer, Lecture Notes in Computer Science 6347, 2010, http://dx.doi.org/10.1007/978-3-642-15781-3_9.

[11] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 3rd ed. edition, 2008, http://dx.doi.org/10.1007/978-3-540-77974-2.

[12] R. Diestel. *Graph Theory, 3rd electronic edition*. Graduate texts in mathematics 173. Springer, 2005.

[13] D. Eppstein. Dynamic generators of topologically embedded graphs. *Proc. 14th ACM-SIAM Symposium on Discrete Algorithms, SODA 2003*, pp. 599–608, 2003, http://dl.acm.org/citation.cfm?id=644108.644208.

[14] J. Fakcharoenphol and S. Rao. Planar graphs, negative weight edges, shortest paths, and near linear time. *J. Comput. Syst. Sci.* 72(5):868–889, 2006, http://dx.doi.org/10.1016/j.jcss.2005.05.007.

[15] J. Fischer and V. Heun. A new succinct representation of rmq-information and improvements in the enhanced suffix array. *Combinatorics, Algorithms, Probabilistic and Experimental Methodologies, First International Symposium, ESCAPE 2007*, pp. 459–470. Springer, Lecture Notes in Computer Science 4614, 2007, http://dx.doi.org/10.1007/978-3-540-74450-4_41.

[16] G. N. Frederickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM J. Comput.* 16:1004–1022, 1987, http://dx.doi.org/10.1137/0216064.

[17] G. N. Frederickson. Planar graph decomposition and all pairs shortest paths. *J. ACM* 38(1):162–204, 1991, http://doi.org/10.1145/102782.102788.

[18] O. Goldreich and D. Ron. Approximating average parameters of graphs. *Random Struct. Algorithms* 32(4):473–493, 2008, http://dx.doi.org/10.1002/rsa.20203.

[19] M. T. Goodrich. Planar separators and parallel polygon triangulation. *J. Comput. Syst. Sci.* 51(3):374–389, 1995, http://dx.doi.org/10.1006/jcss.1995.1076.

[20] T. Husfeldt. Computing graph distances parameterized by treewidth and diameter. *Proc. 11th International Symposium on Parameterized and Exact Computation, IPEC 2016*, pp. 16:1–16:11. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, LIPIcs 63, 2017, http://dx.doi.org/10.4230/LIPIcs.IPEC.2016.16.

[21] P. Indyk. Sublinear time algorithms for metric space problems. *Proc. 31st ACM Symposium on Theory of Computing, STOC 1999*, pp. 428–434. ACM, 1999, http://doi.acm.org/10.1145/301250.301366.

[22] K. Kawarabayashi, P. N. Klein, and C. Sommer. Linear-space approximate distance oracles for planar, bounded-genus and minor-free graphs. *Proc. 38th International Colloquium on Automata, Languages and Programming, ICALP 2011*, pp. 135–146. Springer, Lecture Notes in Computer Science 6755, 2011, http://dx.doi.org/10.1007/978-3-642-22006-7_12.

[23] P. N. Klein. Multiple-source shortest paths in planar graphs. *Proc. 16th ACM-SIAM Symposium on Discrete Algorithms, SODA 2005*, pp. 146–155, 2005, http://dl.acm.org/citation.cfm?id=1070432.1070454.

[24] P. N. Klein, S. Mozes, and C. Sommer. Structured recursive separator decompositions for planar graphs in linear time. *Proc. 45th ACM Symposium on Theory of Computing, STOC 2013*, pp. 505–514, 2013, http://doi.acm.org/10.1145/2488608.2488672. See http://arxiv.org/abs/1208.2223 for the full version.

[25] P. N. Klein, S. Mozes, and O. Weimann. Shortest paths in directed planar graphs with negative lengths: A linear-space $o(n \log^2 n)$-time algorithm. *ACM Trans. Algorithms* 6(2):30:1–30:18, 2010, http://doi.acm.org/10.1145/1721837.1721846.

[26] P. N. Klein and S. Subramanian. A fully dynamic approximation scheme for shortest paths in planar graphs. *Algorithmica* 22(3):235–249, 1998, http://dx.doi.org/10.1007/PL00009223.

[27] R. Klein. *Concrete and Abstract Voronoi Diagrams*. Lecture Notes in Computer Science 400. Springer, 1989, http://dx.doi.org/10.1007/3-540-52055-4.

[28] R. Klein. Abstract voronoi diagrams. *Encyclopedia of Algorithms*, pp. 1–5. Springer Berlin Heidelberg, 2014, http://dx.doi.org/10.1007/978-3-642-27848-8_603-1.

[29] R. Klein, E. Langetepe, and Z. Nilforoushan. Abstract voronoi diagrams revisited. *Comput. Geom.* 42(9):885–902, 2009, http://dx.doi.org/10.1016/j.comgeo.2009.03.002.

[30] R. Klein, K. Mehlhorn, and S. Meiser. Randomized incremental construction of abstract voronoi diagrams. *Comput. Geom.* 3:157–184, 1993, http://dx.doi.org/10.1016/0925-7721(93)90033-3.

[31] R. Lipton, D. Rose, and R. Tarjan. Generalized nested dissection. *SIAM J. Numer. Anal.* 16(2):346–358, 1979, http://dx.doi.org/10.1137/0716027.

[32] D. Marx and M. Pilipczuk. Optimal parameterized algorithms for planar facility location problems using voronoi diagrams. *Algorithms - ESA 2015 - 23rd Annual European Symposium*, pp. 865–877. Springer, Lecture Notes in Computer Science 9294, 2015, http://dx.doi.org/10.1007/978-3-662-48350-3_72. Full version available at http://arxiv.org/abs/1504.05476.

[33] S. Mozes, Y. Nussbaum, and O. Weimann. Faster shortest paths in dense distance graphs, with applications. *CoRR* abs/1404.0977, 2014, http://arxiv.org/abs/1404.0977.

[34] S. Mozes and C. Sommer. Exact distance oracles for planar graphs. *Proc. 23rd ACM-SIAM Symposium on Discrete Algorithms, SODA 2012*, pp. 209–222, 2012, http://dl.acm.org/citation.cfm?id=2095116.2095135.

[35] S. Mozes and C. Wulff-Nilsen. Shortest paths in planar graphs with real lengths in $o(nlog^2/\log\log n)$ time. *Algorithms - ESA 2010, 18th Annual European Symposium, Part II*, pp. 206–217. Springer, Lecture Notes in Computer Science 6347, 2010, http://dx.doi.org/10.1007/978-3-642-15781-3_18.

[36] G. F. I. Y. Nussbaum, P. Sankowski, and C. Wulff-Nilsen. Improved algorithms for min cut and max flow in undirected planar graphs. *Proc. 43rd ACM Symposium on Theory of Computing, STOC 2011*, pp. 313–322, 2011, http://doi.acm.org/10.1145/1993636.1993679.

[37] J. K. Park and C. A. Phillips. Finding minimum-quotient cuts in planar graphs. *Proc. 25th ACM Symposium on Theory of Computing, STOC 1993*, pp. 766–775, 1993, http://doi.acm.org/10.1145/167088.167284.

[38] V. Patel. Determining edge expansion and other connectivity measures of graphs of bounded genus. *SIAM J. Comput.* 42(3):1113–1131, 2013, http://dx.doi.org/10.1137/100811416.

[39] L. Roditty and V. Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. *Proc. 45th ACM Symposium on Theory of Computing, STOC 2013*, pp. 515–524, 2013, http://doi.acm.org/10.1145/2488608.2488673.

[40] J. Snoeyink. Point location. *Handbook of Discrete and Computational Geometry, Second Edition.*, pp. 767–785. Chapman and Hall/CRC, 2004, http://dx.doi.org/10.1201/9781420035315.pt4.

[41] R. Tamassia and G. Liotta. Graph drawing. *Handbook of Discrete and Computational Geometry, Second Edition.*, pp. 1163–1185. Chapman and Hall/CRC, 2004, http://dx.doi.org/10.1201/9781420035315.ch52.

[42] M. Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *J. ACM* 51(6):993–1024, 2004, http://doi.acm.org/10.1145/1039488.1039493.

[43] O. Weimann and R. Yuster. Approximating the diameter of planar graphs in near linear time. *ACM Trans. Algorithms* 12(1):12, 2016, http://doi.acm.org/10.1145/2764910.

[44] C. Wulff-Nilsen. Wiener index, diameter, and stretch factor of a weighted planar graph in subquadratic time. Tech. Rep. 08-16, Department of Computer Science, University of Copenhagen, 2008, http://www.diku.dk/OLD/publikationer/tekniske.rapporter/rapporter/08-16.pdf. Preliminary version in EurCG 2009.

[45] C. Wulff-Nilsen. Wiener index, diameter, and stretch factor of a weighted planar graph in subquadratic time, 2010. Paper M in the PhD thesis of C. Wulff-Nilsen, available at http://www.diku.dk/forskning/phd-studiet/phd/ThesisChristian.pdf.